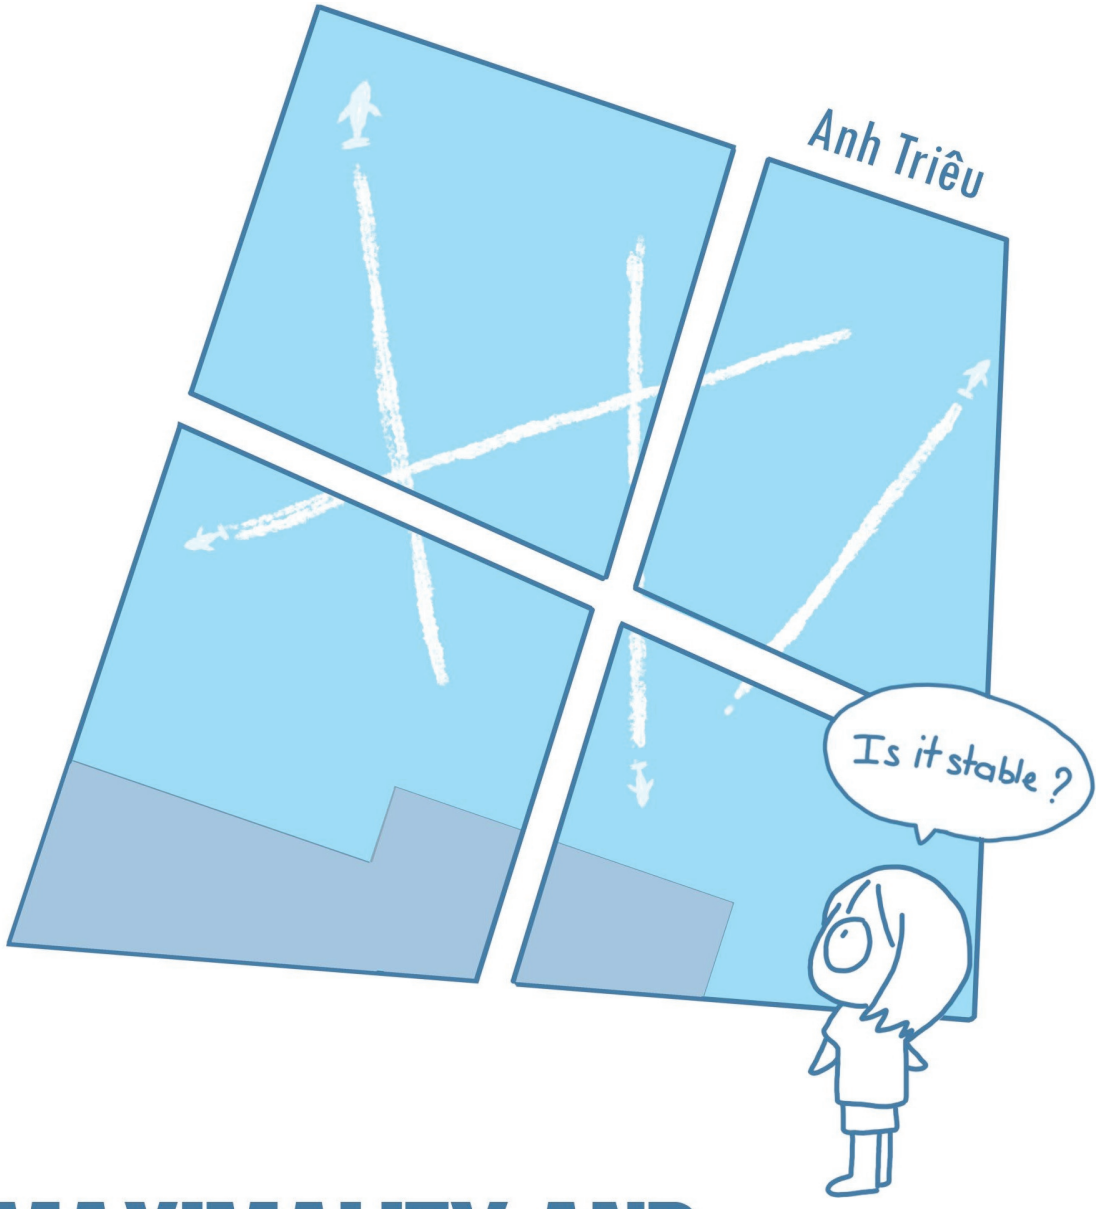

MAXIMALITY AND STABILITY

in matching markets

Propositions accompanying the PhD-thesis

Maximality and Stability in Matching Markets

Anh Triêu, December 7 2026

1. Total utility maximization is often treated as the ”neutral” benchmark in mechanism design, but the choice to maximize aggregate surplus rather than the number of people served is itself a value judgment, not a technical default. (Chapter 2-4)
2. An efficient mechanism can be socially undesirable if society values something other than efficiency. (Chapter 2, 3)
3. The hardest part of mechanism design is not designing the rules, but deciding what the rules should achieve. (Chapter 2, 3)
4. Property rights are valuable precisely because they allow someone to say “no.” (Chapter 5)
5. The fundamental problem of economic information is not that people know different things, but that they have reasons not to tell each other what they know.
6. A model can be mathematically correct and economically wrong.
7. When economic behavior is undesirable, economists should first question the rules of the game before questioning the preferences of the players.
8. A good mathematical proof, like a good joke, has failed if it needs to be explained. *(Inspired by Professor Dries Vermeulen.)*
9. A title should describe what is distinctive, not merely what is true. *(Inspired by Dr. Iwan Bos.)*
10. A good supervisor does not tell you whether your theorem is true; they try to find a counterexample. *(Inspired by Dr. Marc Schröder.)*
11. Good mathematical writing should not require the reader to think recursively. *(Inspired by Professor Bettina Klaus.)*
12. Probabilities do not exist in the world; they exist in our models of the world.

Doctoral thesis

MAXIMALITY AND STABILITY IN MATCHING MARKETS

Anh Triêu

2026

MAXIMALITY AND STABILITY IN MATCHING MARKETS

Dissertation

To obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus, Prof. Dr. Jan Smits,
in accordance with the decision of the Board of Deans,
to be defended in public
on Monday, December 7th 2026, at 13:00 hours

by

Viên Gia Anh Triêu

Supervisor
Prof. Dr. Dries Vermeulen, Maastricht University

Co-supervisors
Dr. Iwan Bos, Maastricht University
Dr. Marc Schröder, Maastricht University

Assessment Committee
Chair:
Prof. Dr. Rudolf Müller, Maastricht University

Members:
Dr. Tom Demeulemeester, Maastricht University
Dr. Andrew Mackenzie, Rutgers University, New Brunswick
Prof. Dr. Britta Peis, RWTH Aachen University

© Anh Triêu, Maastricht 2026.
All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the author.

Cover Anh Triêu, 2026
Production ProefschriftMaken || proefschriftmaken.nl
ISBN 978-94-6534-722-6

To my loved ones,

Acknowledgments

Abstract

This chapter studies the mechanism through which a doctoral dissertation is successfully completed. The analysis considers a repeated interaction between a PhD candidate and a set of agents, including supervisors, colleagues, friends, and family, who provide support in various forms. The main result establishes that completion is not implementable under a mechanism relying solely on the candidate's effort. Instead, sustained encouragement, guidance, and companionship are necessary to satisfy the relevant incentive constraints. The equilibrium characterized in this chapter is unique, efficient, and would not have existed without all of you.

Keywords: *gratitude, support, family, friends, colleagues, thank you.*

Introduction

People often say that success comes from hard work and a little bit of luck. While I never consider myself lucky in the conventional sense, I know that I am incredibly fortunate in the people who have crossed my path. Every important step of my PhD journey has been made possible by the kindness, encouragement, and guidance of the great people I have the privilege to meet.

With Gratitude

First, I would like to thank my three wonderful supervisors. Thank you, *Dries*, for your endless kindness, patience, guidance, and support throughout this journey. I will always remember whenever my proofs were not good enough, you would simply look at me and say, "Still a long way there, Anh." Somehow, those few words always carried both honesty and encouragement. I am also grateful for many small moments we shared outside of research such as dinners at your place, the symposium trip, where you almost missed the train because you were waiting at the wrong platform. They were such small and funny moments, but they captured the human side of this journey so well that I even turned them into a little manga page.

If there is one person who taught me to appreciate counterexamples, it is *Marc*. I am grateful to you for the countless counterexamples, stimulating discussions, and valuable insights throughout my PhD. I appreciate your honesty about my progress and your willingness to point out where I could improve, while always remaining approachable whenever I had questions. Beyond our research discussions, I am grateful for all happy memories we shared, from seeing you at the restaurant where I used to work, to skiing at SnowWorld, and the two bouldering sessions (where, somehow, I always ended up injured when you were around). These moments made this journey much more enjoyable. I am also grateful for the cute Pokémon figure, the cat cookies from Japan, and for introducing me to Karl.[1]

Thank you, *Iwan*, for your guidance and support throughout this journey. I will remember our discussions and many insights you shared with me, which often made me look at problems from a different perspective. I am also grateful for your support and recommendations regarding my lecturer position. Your encouragement and trust mean a lot to me.

Another beautiful coincidence in my journey was meeting *Bettina* at a conference. Thank you for your thoughtful comments on my presentation back then, for welcoming me to Lausanne for three months, and for all the wonderful meals you prepared for me. I am still obsessed with the meringue mango-rhubarb dessert you made – it remains one of my favorite desserts I have ever tasted. I also cherish our hiking walk in Lavaux. Thank you for being so kind, supportive, and caring. Sometimes I wonder what I must have done in my past life to be lucky enough to meet someone like you in this one.

I would also like to thank all members of my reading committee, *Rudolf Müller*, *Britta Peis*, *Andrew Mackenzie*, and *Tom Demeulemeester*, for dedicating their time and effort to evaluating my thesis and for providing insightful comments and suggestions.

I am grateful to all professors and colleagues from the Department of Quantitative Economics, who made the department such a warm and welcoming place to spend the past few years. Research is only one part of a PhD, and so many of my happiest memories came from the people around me.

Andrés, Janos, Chris, and *Anna*, thank you for all of the fun MLSE seminars and insightful discussions. *Bas*, I still laugh whenever I think about the funny moments we shared at both my first and my last conference as a PhD. *Tom*, thank you for the warm welcome in Lausanne, the lovely walk with Bettina through Lavaux, and your thoughtful comments on my dissertation. I am also grateful to our two secretaries *Mandy*, *Dany*, and other professors *Tjark, Ines, Lissa, etc.* for all of your support.

Enrico, thank you for always being there whenever I wanted to share something. I know you could not always relate to my (over)reactions, but you still listened, stayed by my side, and always tried to include me in all of the conversations. I appreciate

[1] Please tell Karl I gratefully acknowledge him for being the best boy cat.

those moments a lot. *Lucas*, you are an amazing office mate. Thank you for putting up with me, my loud keyboard, and all of my boring stories. It is always very fun to have you in the office. *Ivan*, thank you for all the fun Pokedoku sessions. Even though I still have not managed to convince you to watch Frieren, I have no intention of giving up. *Marie*, your kindness and unwavering support mean so much to me. I have always admired not only your brilliant mind, but also your remarkable emotional intelligence. *Zhilin*, my brunch mate! I am so happy our paths crossed. I will always cherish our memories together, especially our trip to Paris and your words to me during our (hopefully not) last dinner in Paris. Thank you *Ashkan* for all of the fun game nights and dinners, and your help when I moved to my new place. It is so fun to randomly run into you whenever I cat-sit for Marc. *(Uncle) Tim*, I will always remember our coffee breaks and the enjoyable walks through Stadpark after lunch. Thank you *Arman* for giving me the opportunity to take care of Tony, another best boy cat. *Ties*, your sense of humor never fails to make me laugh. *Pauline*, thank you for a wonderful trip to Naples and all of the fun lunches we shared. Finally, I would like to thank *Marie Ternes, François, Freija, Farzaneth*, and *Mathijs* for all of the fun game nights and PhD movie nights; *Shasha, Yanru, Xuanzhu, Jing*, and *Ziyi* for all of the amazing meals and all of the cutest gifts you prepared for me; *Shahrezad, Ugur, Lison, Teun, Valentin*, and *Paul* for all the enjoyable lunches, conversations, and moments we shared together. I hope there will be many more to come, at least with those of you who are staying around. My friends are often jealous that I get along so well with my colleagues. The truth is, I have never really thought of you as just colleagues. You all have always been my friends.

Thank you *Pedro* and *Michal* for giving me the opportunity and helping me organize the MLSE seminar. I am also grateful to all professors and colleagues from Department of Microeconomics and Public Economics for welcoming me to the department for my (not so) new position.

Yang, thank you for being such a kind co-author. I am grateful for your patience in putting up with me and my messages at all hours of the day whenever I came up with another idea for our proofs. It has been a pleasure working with you. I am also grateful to everyone from University of Lausanne for making my three-month research visit one of the happiest periods of my PhD. My sincere thanks go to *Jinfeng, Tanzhou, Luxi, Maria, Elena, Geoffrey, Mojgan,* and *Fedor* for making me feel so welcome from the very beginning.

Lina, I'm grateful for all of our dinners, Pokémon Go weekly challenges, and lucky trades with you (and your family members). I hope we can still meet and enjoy our days together.

Thank you *Jane*, *Lily*, and everyone at La Chine for your kindness and support even when I am not working there anymore.

All of this would never happen without *Niels*, who was one of the first people to ask me what I wanted to pursue after my master's degree. Through a fortunate coinci-

dence, he overheard that Dries was looking for a PhD candidate and introduced me to my three wonderful supervisors.

To my two best friends and family,

Cảm ơn hai bạn mình, *Hân* và *Bọt*. Anh không nghĩ mình có thể đi đến ngày hôm nay nếu không có hai người. Anh sẽ luôn nhớ những chuyến đi du lịch siêu vui với Hân và những lần về ngồi tám với Bọt dù hai đứa chẳng ai có drama gì để kể. Cảm ơn vì những cuộc trò chuyện xuyên biên giới, xuyên thời gian. Cảm ơn vì những lời ủng hộ, động viên của cả hai. Tự nhiên viết tiếng Việt lại không biết phải viết gì cho hay nữa. Nói chung là cảm ơn hai người rất, rất nhiều. Hy vọng 10, 20 hay 50 năm nữa tụi mình vẫn sẽ là bạn thân của nhau.

Cuối cùng, con xin cảm ơn *ba mẹ* đã luôn yêu thương con. Con biết để con có được ngày hôm nay, ba mẹ đã phải hy sinh rất nhiều. Cảm ơn ba mẹ vẫn luôn tin tưởng và để con tự theo đuổi những điều mình yêu thích. Không có sự yêu thương và hậu thuẫn của ba mẹ, con đã không thể đi được đến ngày hôm nay. Con hy vọng ba mẹ sẽ luôn tự hào về con, bởi mọi thành quả của con đều có một phần rất lớn là nhờ ba mẹ. Cảm ơn *ngoại, má, các cậu, các mợ, anh trai, Yến* và *Nhi* vì đã luôn ở bên con/Siu. *Bé Vy*, cảm ơn em bé vì những lần lau kệ sách cho Siu, làm Siu cười, và những drama học đường vô cùng hấp dẫn. *Chị Thư*, cảm ơn chị vì đã luôn bên cạnh động viên và ủng hộ em.

Conclusion

To conclude, the main theorem of this chapter is simple: no dissertation is written by a single person. The proof spans four years and relies on countless conversations, words of encouragement, patient advice, and moments of support. Although this chapter can only approximate the extent of my gratitude, I hope it provides a meaningful lower bound. Future work consists of maintaining these relationships well beyond the completion of this thesis, with no planned end date.

Conflict of Interest Declaration: The author declares a strong bias towards everyone mentioned in this chapter.

Contents

Acknowledgments **vii**

1 Introduction **1**
1.1 Matching Markets with Transfers . . . 3
1.2 Matching Markets with Endowment Structures . . . 5

2 Maximal Matchings **7**
2.1 Introduction . . . 8
2.2 Model . . . 11
2.3 The Flip Algorithm . . . 13
2.4 The Price of Matching Maximization . . . 18
2.5 Concluding Remarks . . . 21
Appendix A: A Justification for Maximal Matchings . . . 23
Appendix B: Proofs . . . 24

3 Monopolistic Trade Maximization **33**
3.1 Introduction . . . 34
3.2 Model . . . 35
3.3 Total Utility Maximization as Upper Bound . . . 38
3.4 Average Allocation Rules . . . 39
3.5 Bilateral Value-Creation . . . 43
3.6 Multi-Unit Demand . . . 44
3.7 Concluding Remarks . . . 45
Appendix: Proofs . . . 46

4 Total Utility Maximization with Heterogeneous Valuations **55**
4.1 Introduction . . . 56
4.2 Preliminaries . . . 57
4.3 A TU-Maximizer Characterization . . . 61
4.4 Concluding Remarks . . . 70
Appendix: Proofs . . . 72

5 Serial Top Trading Cycles and the Exclusion Core **75**
5.1 Introduction . . . 76

5.2 Preliminaries 79
5.3 Serial Top Trading Cycles 86
5.4 Main Results 92
5.5 Concluding Remarks 101
Appendix A: Proofs of the results in Section 5.2 102
Appendix B: Proofs of the results in Section 5.3 103
Appendix C: Proofs of the results in Section 5.4 111
Appendix D: An illustration of the STTC algorithm and stability results . . 117
Appendix E: Acyclic endowment structures 123

Bibliography **133**

Valorisation **139**

Summary **143**

Samenvatting **147**

About the author **151**

1

Introduction

Many economic interactions are like a game of poker where everyone is holding their cards close to their chest. That is, agents typically possess information about their own valuations, costs, or preferences that is not directly observable to others. This is called private information. For example, on the housing market, buyers know how much they value a property, while sellers know their opportunity costs. When bidding for a procurement contract, construction firms have private information about their construction costs. When choosing a school, students privately rank schools according to their preferences. In each of these settings, outcomes depend not only on preferences and feasibility constraints, but also on the strategic use of private information.

When information is asymmetric, agents acting in their own interest need not generate socially desirable outcomes. Buyers may understate their willingness to pay, sellers may inflate reservation prices, and participants may misrepresent preferences if doing so improves their payoff. As a result, mutually beneficial trades may fail to occur, allocations may be inefficient, and stable arrangements may break down. The famous Gibbard-Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975) shows that in almost any fair voting system, there is an incentive for people to "vote strategically". For example, a voter may be better off voting for a candidate they don't actually like just to stop an even more disliked candidate from winning. These observations highlight the importance of designing institutions that account for individual incentives and informational constraints.

Mechanism design provides a framework to address these challenges. This branch of economics and game theory studies how individual incentives affect collective results, and how well-designed rules can align personal behavior with social goals. It is often described as "reverse game theory". Unlike traditional game theory which takes the rules of the game as given and analyzes agents' behavior under those rules, mechanism design adopts the perspective of a third party, such as a policymaker, regulator, or platform designer, who begins with an objective and then designs the rules of interaction to achieve it. Assuming that agents possess private information and act strategically, the designer has to construct mechanisms that induce desirable outcomes despite informational asymmetry or individual incentives. The importance of the field is reflected in the Nobel Memorial Prize in Economic Sciences awarded to Leonid Hurwicz, Eric Maskin, and Roger Myerson for their foundational contributions to mechanism design theory.

The foundations of mechanism design can be traced back to earlier debates on the role of information in economic systems. In particular, the inconclusive discussions on socialism versus capitalism highlighted the limits of classical price theory in environments where information is private. As emphasized by Hayek (1945), the central economic problem is not the allocation of given resources, but the effective use of dispersed information held by different individuals. Building on this insight, Hurwicz (1972) introduced the notion of incentive compatibility and formalized mechanisms as rules that map agents' private information into collective decisions. Subsequent developments, including the Bayesian framework of Harsanyi (1968) and the revelation principle (Myerson, 1979, 1981), established mechanism design as a theory of how institutions can be constructed to aggregate private information and align individual incentives with social objectives.

Building on these theoretical foundations, the literature has developed at the intersection of auction theory, social choice theory, and contract theory. Its insights have been applied in a wide range of settings. In school choice, mechanisms determine how students are assigned to schools based on reported preferences (see, *e.g.*, Abdulkadiroğlu & Sönmez, 2003; Abdulkadiroğlu et al., 2005; Gale & Shapley, 1962; Haeringer & Klijn, 2009). In kidney exchange programs, allocation rules coordinate compatible donor–patient pairs to maximize the number of successful transplants (see, *e.g.*, Roth et al., 2004, 2005a, 2005b). In auctions, well-designed rules ensure efficient allocation of goods while providing incentives for truthful reports (see, *e.g.*, Clarke, 1971; Groves, 1973; Myerson, 1981; Vickrey, 1961). Similarly, online marketplaces rely on pricing and matching mechanisms to coordinate transactions among agents with private information (see, *e.g.*, Deng et al., 2010; Gallien, 2006).

This thesis applies mechanism design to two different types of matching markets that illustrate different strategic and/or institutional challenges. The first part of the thesis analyzes matching markets where transfers are allowed. Typical examples include housing markets, labor markets, procurement auctions, and online marketplaces. In such settings, the central challenge is to design mechanisms that generate value-creating

trades and achieves some social objectives, while ensuring that agents truthfully reveal their private information.

The second part turns to matching markets with an endowment structure, where initial ownership rights determine feasible allocations. In these markets, agents possess initial ownership rights over indivisible goods, and these rights constrain which reallocations are feasible. Trade is not only governed by agents' preferences, but also by property rights: an object cannot be reassigned without the consent of its owner. Typical examples include housing markets with existing tenants (Abdulkadiroğlu & Sönmez, 1999), school seat reallocation when certain seats are reserved (see, *e.g.*, Hafalir et al., 2013), or allocation of offices within a university where different faculties initially occupy specific offices. In these environments, the central concern is not only efficiency, but also stability with respect to ownership rights. An allocation must ensure that no group of agents can use their rights to reassign objects among themselves in a way that makes them all strictly better off.

Studying both environments is important because they capture different dimensions of real-world allocation problems. Markets with transfers emphasize the role of allocations and prices in overcoming private information. Markets with endowments emphasize the role of property rights and the constraints they impose on feasible reallocations. While both settings involve strategic behavior, the appropriate design principles differ. Together, they provide a more complete understanding of how matching mechanisms function under different institutional constraints.

1.1 Matching Markets with Transfers

The first part of the thesis studies matching markets in which certain monetary transfers are allowed and feasibility is determined by demand and supply constraints rather than by agents' initial ownership rights. The central question is how to design mechanisms that determine who trades with whom and at what payments given some objective, while accounting for private information and strategic behavior.

To illustrate, consider a housing market where consumers have cardinal valuations (maximum willingness to pay) for houses and property developers incur (opportunity) costs of providing a house to the market. These valuations and costs are typically private information. A trade between a developer and a consumer is value-creating when the consumer's valuation exceeds the property developer's cost, so that a positive surplus can be generated. However, because valuations and costs are privately known, consumers and property developers may have incentives to misreport their information in order to obtain a more favorable price. If the mechanism used to organize trade does not properly account for these incentives, some value-creating transactions may fail to occur, inefficient trades may be implemented, and consequently the social objective may not be achieved.

1.1.1 Maximality

In this thesis, we consider situations in which the social objective is to maximize the number of value-creating matches, which we refer to as *maximal matchings*. Among allocation rules that aim to maximize the number of value-creating matches, we focus on those that satisfy two additional properties. First, we require that an allocation rule selects at least all trades that would arise under a total-utility maximizing allocation, and possibly more whenever these additional matches are value-creating. Second, we require the mechanism to satisfy *incentive compatibility*, ensuring that agents truthfully report their private information. The following subsections introduce these two concepts in turn.

Comparison with total-utility maximization

In the mechanism design literature, a total-utility maximizing allocation rule (TU-maximizer) is the standard social objective. Maximizing total utility corresponds to the surplus-maximizing allocation in the assignment game and underlies competitive equilibrium outcomes (L. S. Shapley & Shubik, 1971). Furthermore, such allocations are Pareto optimal, meaning no feasible matching can make one agent better off without making another worse off. While such mechanisms are efficient in the utilitarian sense, they may ignore potential value-creating matches.

We therefore focus on allocation rules that improve upon a TU-maximizer in terms of the number of value-creating matches. The objective is to increase the number of transactions relative to TU-maximization, while ensuring that all transactions remain value-creating. In this sense, the rule weakly increases the number of matches relative to a TU-maximizer, while avoiding the inefficiencies that may arise from maximizing the number of trades alone.

Incentive compatibility

Since valuations and costs are private information, the social planner applies a direct mechanism in which agents are asked to report their private information. Then, allocations and payments are determined as a function of these reports.

Incentive compatibility ensures that agents find it optimal to reveal their true valuations and costs despite the presence of strategic incentives. That is, a mechanism is incentive compatible if truthful reporting maximizes each agent's utility given the mechanism's allocation and payment rules. In other words, for every agent, the utility obtained from reporting their true valuations or preferences is at least as high as the utility obtained from reporting any other values.

1.1.2 Overview of Chapters 2–4

Chapter 2 studies maximal matchings in a unit-demand, unit-supply environment with homogeneous goods in which each consumer demands at most one indivisible good and each producer supplies a single unit at a cost. Chapter 3 extends these concepts to a unit-demand setting, where consumers still demand at most one unit, and a single producer supplies several identical units to the market.

Chapter 4 turns to a unit-demand environment with heterogeneous goods, in which consumers have different valuations across objects. In contrast to the previous chapters, there are no longer two strategic sides, as producers' costs are normalized to zero and assumed to be common knowledge. Strategic behavior arises solely on the demand side. This specification isolates the allocation problem driven by heterogeneity across goods, while preserving the incentive constraints faced by consumers.

1.2 Matching Markets with Endowment Structures

The second part of the thesis studies matching markets in which agents possess initial ownership rights over indivisible goods. In these settings, feasibility and participation constraints are determined by agents' endowments, and desirable allocations are evaluated using stability concepts. The central question is how to design allocation rules that respect ownership rights, prevent deviations by agents, and achieve outcomes that are stable given the initial endowment structure.

To illustrate, consider a housing market in which agents may privately own one or more houses or may own none at all. Houses may be privately owned by individual agents or collectively owned by all or a group of agents. Moreover, some houses may already be occupied or assigned to agents who are different from their legal owners. Agents have ordinal preferences for houses. In such settings, reallocating houses is not simply a matter of matching preferences with availability. Any reassignment must respect the structure of ownership, and groups of agents may be able to reorganize allocations by exercising their property rights. These features shape which allocations are stable and which are vulnerable to deviations.

1.2.1 Exclusion Core Stability

Traditional stability concepts view endowments as bundles of goods that agents can trade or consume. This interpretation works well in simple ownership environments, but it becomes problematic when the endowment structure is more complex, for example, when objects are jointly owned by several agents. In such cases, some traditional stability notions rule out allocations that seem feasible given the ownership structure, while others allow allocations that are difficult to justify from a property-rights perspective. Balbuzanov and Kotowski (2019) propose a different interpretation. They view endowments as systems of exclusion rights: owning an object means

having the right to exclude others from using it. From this perspective, feasible deviations are those that can be implemented by exercising these exclusion rights. For example, suppose an agent owns an object. If an allocation assigns the object to another agent and the owner prefers their own object to the currently assigned one, they can block the allocation and reclaim their own object.

Building on this idea, they introduce the notion of the *exclusion core*, a stability concept that restricts blocking to coalitions whose deviations are feasible through exclusion. An allocation is exclusion core stable if no coalition of agents can use their exclusion rights to reassign objects among themselves such that every member of the coalition is strictly better off.

1.2.2 Overview of Chapter 5

Chapter 5 studies the concept of exclusion core stability in matching markets with mixed-ownership. Given a priority order over agents, an endowment structure, an initial allocation, and agents' preferences, we construct directed graphs, referred to as serial top trading cycles (STTC) graphs. These STTC graphs, together with the associated STTC algorithm, turn out to be a key tool for analyzing exclusion core stable allocations.

2

Maximal Matchings

This chapter is based on Triêu, A., Bos, I., Schröder, M., & Vermeulen, D. (2026). Maximal matchings. *Journal of Mathematical Economics*, *125*, 103254. https://doi.org/10.1016/j.jmateco.2026.103254

Earlier versions of this work were presented at the 2024 Coalition Theory Network Workshop, the 2024 Oligo Workshop, the 7th world congress of the Game Theory Society (GAMES, 2024), the MLSE at Maastricht University and the Microeconomics Seminar at University of Lausanne.

"It is the greatest happiness of the greatest number that is the measure of right and wrong."
(Bentham, 1776)

2.1 Introduction

Economics is often described as a social science that studies the production, distribution, and consumption of wealth.[1] Throughout its history, a good number of criteria have been developed to judge levels and distributions of welfare. Perhaps the most popular one is the so-called utilitarian rule. Roughly speaking, this rule says that goods and services should be produced and allocated in such a way that the combined utility of all agents involved is maximal.

The utilitarian rule has been criticized on many grounds, with the most common one being its neglect regarding issues of distributive justice.[2] Indeed, abiding by this principle may result in situations where most, if not all, wealth ends up in the hands of a few.[3] This can be particularly problematic when the goods or services are considered 'essential' such as electricity, food, housing, medical supplies, *et cetera*. In such cases, one can imagine policymakers to be more concerned with the availability and accessibility of wealth rather than with its aggregate value. That is to say, it may be preferred to have supplies produced and allocated in such a way that it serves as many parties as possible.[4]

In this chapter, we suppose that the social goal is to maximize the number of buyer-seller matches, which we coin a *maximal matching*.[5] We study maximal matchings in a setting that is comparable to Böhm-Bawerk's classic horse market game.[6] There are multiple (potential) sellers that can offer one unit of an indivisible, homogeneous product and multiple (potential) buyers that consume at most one unit of the good. This setting fairly well describes the market for tract housing, for example.[7] Note that competition is two-sided and that interaction between consumers and producers takes

[1] Definitions along these lines date back to *A Treatise on Political Economy; or the Production, Distribution, and Consumption of Wealth* by Say (1836). For a detailed discussion, see Backhouse and Medema (2009).

[2] See, *e.g.*, Feldman (1974), Mandler (1999) and, more recently, Graafland (2021).

[3] As vividly formulated by Nozick (1974, p. 41): "*Utilitarian theory is embarrassed by the possibility of utility monsters who get enormously greater gains in utility from any sacrifice of others than these others lose. For, unacceptably, the theory seems to require that we all be sacrificed in the monster's maw, in order to increase total utility.*"

[4] It is further worth noting that utilitarian (re-)allocation may require transfers across all market participants. Such compensation schemes require coordination among those participants, which may not be feasible due to market imperfections caused by, for example, spatial or temporal frictions.

[5] It is worth emphasizing that the maximal matching goal is not only applicable within the context of essential goods provision. In digital markets, for example, a platform owner may be primarily interested in maximizing the number of transactions or 'clicks' as this gives valuable information that can be exploited elsewhere. See, *e.g.*, Edelman et al. (2007) and Varian (2007) for an analysis of sponsored search auctions.

[6] Böhm-Bawerk (2012). For an analysis and discussion of (variations of) the horse market game, see, *e.g.*, L. S. Shapley and Shubik (1971) and Núñez and Rafels (2005).

[7] Tract houses are houses with a standardized design and sometimes referred to as 'cookie cutter houses'. For a detailed analysis and discussion, see Mackenzie and Zhou (2022) as well as some of the references therein.

place only through bilateral transactions, *i.e.*, deals made directly between a buyer and a seller. This is characteristic for many markets, including the housing market, the medical supplies market, the private transportation market (*e.g.*, taxi services), and the recently emerging peer-to-peer electricity markets. It is worth emphasizing that the bilateral nature of market deals requires matches to be value-creating, *i.e.*, the valuation of the buyer must exceed the cost of the seller.[8]

To fix ideas, consider the following example of two families looking for a house. Suppose there are two property developers, each of whom has an apartment available for sale. The apartments are virtually identical, but the costs of selling the place differ. After attending the open house, both families learn their valuation for the accommodation. Table 1 gives an overview of the valuations (v) and costs (c), which can be thought of as reservation prices. Family 1 values either apartment at 10, whereas family 2 values both at 5. The developers' costs are 2 and 6, respectively.

v_1	v_2
10	5
c_1	c_2
2	6

Table 2.1: A market with 2 consumers and 2 producers.

Note that if the policymaker's goal is to maximize total utility, it would prefer the competitive equilibrium in which project developer 1 sells to family 1. This, however, leaves family 2 empty-handed, because project developer 2's minimum price ($c_2 = 6$) exceeds this family's valuation ($v_2 = 5$). In this scenario, total utility is $10 - 2 = 8$. Yet, if the policymaker's goal is to have housing for as many people as possible, then it would like project developer 1 to trade with family 2 and project developer 2 to sell to family 1. This yields a total utility of $(5 - 2) + (10 - 6) = 7$ and a maximal number of matches, namely 2.[9]

What makes such a maximal matching problem challenging is that costs and valuations are private information. Following the above example, a family's actual valuation for a house is unknown to the project developer as well as to competing buyers. Similarly, a project developer often has a pretty good idea about the (opportunity) cost of selling the property, but that information is typically not available to any of the other market players. Moreover, both buyers and sellers are likely to pursue goals other than maximizing the total number of deals between them. This raises the question

[8] Note that by focusing on value-creating bilateral transactions, we exclude the possibility of a social planner ordering goods from producers to reallocate them to consumers.

[9] Note that in a competitive equilibrium, there is a uniform price for all buyer-seller matches. By contrast, for maximal matchings, we need personalized prices: the price for the first pair (project developer 1 and family 2) should be between 2 and 5 and the price for the second pair (project developer 2 and family 1) should be between 6 and 10. We thank an anonymous referee for pointing this out.

of whether one can design a mechanism that achieves a maximal matching. In what follows, we show that the answer to this question is affirmative.

Specifically, we propose a mechanism with the following properties. First, all agents involved are incentivized to reveal their true valuations or costs, which implies that the mechanism is implementable. Second, the mechanism identifies the optimal maximal matching in the sense that it selects the maximal matching with the highest total utility. Third, we show that matching maximization yields at most twice the number of transactions in comparison to total utility maximization. Fourth, and similar to the well-known total-utility-maximizing VCG-mechanism, the mechanism is not budget balanced.[10] A policymaker thus has to look for outside funds when it aims for a maximal matching.

By maximizing the number of buyer-seller matches, the proposed mechanism achieves the social goal 'directly'. In principle, a policymaker can achieve a comparable outcome also 'indirectly' through an appropriate (re-)distribution of agents' (initial) endowments. This, however, may prove problematic for at least two reasons. First, the information required to determine the preferred allocation is typically lacking and cannot be elicited truthfully. Second, redistribution methods (*e.g.*, taxes and subsidies) are in themselves inefficient. Apart from this, experimental research suggests free market competition to converge to a competitive equilibrium under plausible conditions, which commonly implies a less than maximal matching.[11] To achieve the latter then requires a central authority to provide the appropriate incentives. It should, moreover, be mentioned that indirect routes to maximal matchings may simply not be available due to technical limitations. For example, storing energy still proves challenging in many situations so that energy suppliers may have to serve their customers instantly and directly.

This research builds on the classic contribution by Myerson and Satterthwaite (1983). They established an 'impossibility result' for the case of a single bilateral trade (*i.e.*, one buyer and one seller) by showing that there exists no efficient, strategy-proof mechanism that is budget balanced. More generally, the mechanism design literature is rich in showing that it is typically impossible to develop a mechanism that meets all desirable properties. Consequently, work in this field focuses on the design of mechanisms that are strategy-proof (by ensuring that the corresponding algorithms satisfy the appropriate monotonicity constraint), while relaxing some of the other properties (*e.g.*, efficiency or balanced budgets). This chapter further relates to work on (algorithmic) mechanism design, the origins of which date back to Nisan and Ronen (1999) and Lehmann et al. (2002).[12] In the context of double auctions, scholars have considered situations in which the market maker pursues objectives other than maximizing the number of bilateral matches, such as revenue maximization (*e.g.*, Deshmukh et al. (2002), Deng et al. (2014)) or budget balancedness (*e.g.*, Colini-Baldeschi et al. (n.d.)).

[10] VCG-mechanisms are named after the contributions of Vickrey (1961), Clarke (1971) and Groves (1973).
[11] See, *e.g.*, Smith (1962), Gode and Sunder (1993), Bosch-Domenech and Sunder (2000), and Lin et al. (2020).
[12] For a survey of this literature, see Nisan (2015).

Our analysis also touches upon the classic efficiency-equity trade-off. While the utilitarian approach yields an efficient allocation of goods or services, matching maximization carries an egalitarian flavor. Indeed, as we show formally in Appendix A, the proposed mechanism with an equal split of surpluses gives an outcome similar to what one would obtain by applying the egalitarian rule. This research is, therefore, naturally connected to the growing body of recent work that takes a mechanism design approach toward social objectives other than total welfare maximization. For instance, Condorelli (2013) argues that market mechanisms may indeed be optimal when a designer pursues allocation objectives not directly related to welfare (*e.g.*, maximizing the number of lives saved). In the same spirit, Dworczak et al. (2021), Groh and Reuter (2023) and Akbarpour et al. (2024) show how market design can be used to reach redistributive objectives when there is heterogeneity in marginal values for money, whereas Finster et al. (2024) propose a mechanism to obtain equality in buyers' surplus in multi-unit auctions.

The chapter is part of a broader inquiry into how far implementable mechanisms can push trade volume beyond the total-utility-maximizing benchmark. Chapter 3, considers this question under a different market structure: a monopolist facing multiple buyers, and shows that the scope for improvement depends critically on the cost structure. Chapter 4 generalizes the question further under a market with heterogeneous products, where agents may value distinct objects differently.

The remainder of this chapter is organized as follows. The next section introduces the modeling framework. Section 2.3 presents the mechanism. Section 2.4 examines the costs of maximal matchings. Section 2.5 concludes. The link to egalitarianism is provided in Appendix A. All proofs are relegated to Appendix B.

2.2 Model

Consider a market comprising a set of $M = \{1, \ldots, m\}$ consumers and a set of $N = \{1, \ldots, n\}$ producers. Each consumer $i \in M$ attempts to purchase one unit of a good, which it values at $v_i \geqslant 0$. Each producer $j \in N$ can produce one unit of this good at a cost of $c_j \geqslant 0$. We assume that valuations and costs are private information. In the ensuing analysis, we use the subscripts i and j to describe consumers and producers, respectively.

Let $T \subseteq M \times N$ be a *matching*, where for every $i \in M$, $|\{j|(i,j) \in T\}| \leqslant 1$, and for every $j \in N$, $|\{i|(i,j) \in T\}| \leqslant 1$. The group of buyers and sellers that are involved in a transaction are indicated by the sets $K \subseteq M$ and $L \subseteq N$, respectively. The total utility generated by a matching T is, therefore, given by:

$$TU(T) = \sum_{i \in K} v_i - \sum_{j \in L} c_j.$$

In what follows, we focus on matches where the valuation of the buyer exceeds the cost of the seller. To that end, we introduce the following two concepts:

Definition 2.2.1 *A match $(i, j) \in T$ is **value-creating** if $v_i \geqslant c_j$.*

Definition 2.2.2 *A matching T is **bilaterally rational** if all $(i, j) \in T$ are value-creating.*

It should be noted that, without transfers, Pareto improvements are impossible in the markets that we consider since each match increases the utility of the consumer, but makes the producer worse off. Welfare gains can, however, be obtained by allowing for bilateral transfers (*i.e.*, transfers from one consumer to one producer). Note that a welfare improving bilateral transfer t with $c_j \leqslant t \leqslant v_i$ exists only when the match is value-creating.[13] One can then define a Pareto frontier with respect to value-creating matches. This frontier consists of matchings for which it holds that there are no unmatched pairs that are value-creating and, in particular, also contains matchings that maximize total utility. We consider a social planner that aims for a matching with a maximal number of matches on this Pareto frontier.

Since information is incomplete, the social planner employs a direct mechanism to reach its goal.[14] To that end, it asks all consumers and producers to report their valuations and costs.[15] Let $\mathbf{r} = (r_1, r_2, \ldots, r_m)$ and $\mathbf{s} = (s_1, s_2, \ldots, s_n)$ be the vectors of reported valuations and costs, respectively. The vector $\mathbf{r}_{-i}$ indicates the reported valuations of all consumers other than i and the vector $\mathbf{s}_{-j}$ denotes the reported costs of all producers other than j. Let R and S be the corresponding sets of all possible reported valuation and cost profiles, respectively. Generic elements are represented by $\mathbf{r}, \mathbf{t} \in R$ and $\mathbf{s}, \mathbf{u} \in S$.

A direct mechanism, $\mathcal{M} = (A, p)$, consists of an *allocation rule* A and a *payment rule* p. An allocation rule $A : (\mathbf{r}, \mathbf{s}) \mapsto A(\mathbf{r}, \mathbf{s}) = (K, L)$ takes the profiles of communicated valuations and costs as input and produces output $A(\mathbf{r}, \mathbf{s}) = (K, L)$. Note that an allocation rule only indicates which buyers (K) and sellers (L) are involved in a trade, whereas a matching additionally specifies the buyer-seller pairs. Each allocation then brings about the following outcome:

$$a_i(\mathbf{r}, \mathbf{s}) = \begin{cases} 1 \text{ if } i \in K \\ 0 \text{ if } i \notin K, \end{cases}$$

[13] Observe that bilateral rationality is a weaker condition than the condition for a total-utility-maximizing matching, which requires a general compensation scheme across all market participants.

[14] If information would be complete, then it is straightforward to select a mechanism that serves the social planner's goal. The adjective 'direct' refers to the revelation principle (see, *e.g.*, Gibbard (1973)), which states that for any mechanism equilibrium there exists an equivalent incentive-compatible direct revelation mechanism.

[15] This is a common approach in mechanism design. By contrast, Peters and Severinov (2006) consider a 'decentralized' dynamic mechanism within a unit-demand, unit-supply setting.

and

$$b_j(\mathbf{r}, \mathbf{s}) = \begin{cases} 1 \text{ if } j \in L \\ 0 \text{ if } j \notin L. \end{cases}$$

The *payment rule* also takes all reported valuations and costs as an input and uses this information to specify the amount $p_i(\mathbf{r}, \mathbf{s})$ that consumer i has to pay and the amount $w_j(\mathbf{r}, \mathbf{s})$ that producer j receives.

2.3 The Flip Algorithm

In this section, we propose an algorithm (coined the 'Flip Algorithm') to identify a matching with a maximum number of value-creating matches. We show that this algorithm is implementable (Section 2.3.1) and that it yields a maximal matching with the highest total utility (Section 2.3.2).

To begin, let us formally describe the Flip Algorithm.

Input: Reported vectors $\mathbf{r}$ and $\mathbf{s}$. Without loss of generality, assume that $r_1 \geqslant r_2 \geqslant \ldots \geqslant r_m$ and $s_1 \leqslant s_2 \leqslant \ldots \leqslant s_n$.

Step 1. Let $T_F = \varnothing$. If $r_1 < s_1$, output T_F and terminate. Otherwise, set $k = 2$ and move to Step k.

Step $k \geqslant 2$. If $r_i \geqslant s_{k+1-i}$ for all $i \in \{1, \ldots, k\}$, set $k = k + 1$ and move to Step k. Otherwise, output $T_F = \{(i, k-i) \mid 1 \leqslant i \leqslant k-1\}$ and terminate.

Observe that if all consumers and producers report their valuation or cost truthfully, then the outcome of the Flip Algorithm is bilaterally rational. It is also noteworthy that the outcome of the Flip Algorithm can be computed in polynomial time. The next example gives an illustration.

Example 2.3.1 *Consider a market with 4 consumers and 4 producers, all of which are assumed to report their true valuation and cost (the issue of 'truth-telling' is taken up in Section 2.3.1 below). Table 2.2 presents the reported valuations and costs.*

r_1	r_2	r_3	r_4
9	8	5	3
s_1	s_2	s_3	s_4
4	6	7	11

Table 2.2: A market with 4 consumers and 4 producers.

Starting with $k = 1$, we compare r_1 to s_1. Since $r_1 \geqslant s_1$, we can proceed to $k = 2$. We then compare r_1 to s_2 and r_2 to s_1. Note that we effectively reverse the order of the first two producers (hence the name Flip Algorithm) to obtain Table 2.3 below. Since $r_1 \geqslant s_2$ and $r_2 \geqslant s_1$, we proceed to $k = 3$. Then, we reverse the order of the first three producers, which gives the situation as presented in Table 2.4. Also in this case, the first three pairs are value-creating so that we proceed to $k = 4$. This results in the situation as presented in Table 2.5. Note that with $k = 4$, valuations fall short of costs for the first, the third and the fourth pair. The Flip Algorithm, therefore, yields the matching $T_F = \{(1, 3), (2, 2), (3, 1)\}$.

r_1	r_2	r_3	r_4
9	8	5	3
s_2	s_1	s_3	s_4
6	4	7	11

Table 2.3: Valuations and costs when $k = 2$.

r_1	r_2	r_3	r_4
9	8	5	3
s_3	s_2	s_1	s_4
7	6	4	11

Table 2.4: Valuations and costs when $k = 3$.

r_1	r_2	r_3	r_4
9	8	5	3
s_4	s_3	s_2	s_1
11	7	6	4

Table 2.5: Valuations and costs when $k = 4$.

2.3.1 Implementability

Let us now address the issue of implementability. Specifically, we examine whether it is a dominant strategy for all consumers and producers to report their valuation or cost truthfully. To that end, we introduce the following three concepts:

Definition 2.3.2 *A mechanism is **dominant-strategy incentive compatible (DSIC)** if, for all $i \in M$, $r_i, t_i \in \mathbb{R}_+$ and $\mathbf{r}_{-i} \in \mathbb{R}_+^{m-1}$, it holds that:*

$$r_i \cdot a_i(r_i, \mathbf{r}_{-i}, \mathbf{s}) - p_i(r_i, \mathbf{r}_{-i}, \mathbf{s}) \geqslant r_i \cdot a_i(t_i, \mathbf{r}_{-i}, \mathbf{s}) - p_i(t_i, \mathbf{r}_{-i}, \mathbf{s}),$$

and, for all $j \in N$, $s_j, u_j \in \mathbb{R}_+$ and $\mathbf{s}_{-j} \in \mathbb{R}_+^{n-1}$, it holds that:

$$w_j(\mathbf{r}, s_j, \mathbf{s}_{-j}) - s_j \cdot b_j(\mathbf{r}, s_j, \mathbf{s}_{-j}) \geqslant w_j(\mathbf{r}, u_j, \mathbf{s}_{-j}) - s_j \cdot b_j(\mathbf{r}, u_j, \mathbf{s}_{-j}).$$

If a mechanism is DSIC, then the net gain of reporting the true valuation or cost is at least as large as the net gain of reporting anything else. This makes reporting truthfully a dominant strategy for both buyers and sellers. Using this concept, implementability is then defined as follows:

Definition 2.3.3 *An allocation rule A is **implementable** if there exists a payment rule p such that (A, p) is DSIC.*

Finally, an allocation rule is monotonic when a consumer does not receive less when it would raise its reported valuation and a producer does not sell less when it would reduce its reported cost. Formally, it is defined as follows:

Definition 2.3.4 *An allocation rule A is **monotonic** if for all $i \in M$, $r_i, t_i \in \mathbb{R}_+$, $\mathbf{r}_{-i} \in \mathbb{R}_+^{m-1}$, and $\mathbf{s} \in \mathbb{R}_+^n$, it holds that: $r_i \leqslant t_i$ if $a_i(r_i, \mathbf{r}_{-i}, \mathbf{s}) = 0$ and $a_i(t_i, \mathbf{r}_{-i}, \mathbf{s}) = 1$, and, for all $j \in N$, $s_j, u_j \in \mathbb{R}_+$, $\mathbf{s}_{-j} \in \mathbb{R}_+^{n-1}$, and $\mathbf{r} \in \mathbb{R}_+^m$, it holds that: $s_j \geqslant u_j$ if $b_j(\mathbf{r}, s_j, \mathbf{s}_{-j}) = 0$ and $b_j(\mathbf{r}, u_j, \mathbf{s}_{-j}) = 1$.*

We now have all the ingredients available to (re-)state Myerson's Lemma.

Lemma 2.3.5 (Myerson's Lemma (1981)) *An allocation rule A is implementable if, and only if, A is monotonic. Moreover, if A is monotonic, then*

$$p_i(r_i, \mathbf{r}_{-i}, \mathbf{s}) = \begin{cases} 0 & \text{if } a_i(r_i, \mathbf{r}_{-i}, \mathbf{s}) = 0, \\ \inf\{t_i \mid a_i(t_i, \mathbf{r}_{-i}, \mathbf{s}) = 1\} & \text{if } a_i(r_i, \mathbf{r}_{-i}, \mathbf{s}) = 1, \end{cases} \text{ for all } i \in M$$

and

$$w_j(\mathbf{r}, s_j, \mathbf{s}_{-j}) = \begin{cases} 0 & \text{if } b_j(\mathbf{r}, s_j, \mathbf{s}_{-j}) = 0, \\ \sup\{u_j \mid b_j(\mathbf{r}, u_j, \mathbf{s}_{-j}) = 1\} & \text{if } b_j(\mathbf{r}, s_j, \mathbf{s}_{-j}) = 1, \end{cases} \text{ for all } j \in N$$

is the unique payment rule (up to an additive constant) such that the mechanism (A, p) is DSIC.

Intuitively, each agent pays or receives its 'switch point' where it goes from 'not being matched' to 'being matched'.

Our first main result shows that the allocation rule induced by the Flip Algorithm, denoted by A_F, is monotonic.

Theorem 2.3.6 *The allocation rule A_F is monotonic.*

We refer to the mechanism induced by the Flip Algorithm as the *Flip Mechanism* (A_F, p_F). The following example illustrates the payment rule for the market as described in Example 2.3.1.

Example 2.3.7 *We compute the price for consumer $i = 1$ in Example 2.3.1. If consumer 1 reports a valuation below 5, then it is unmatched by the Flip Algorithm as the number of matches decreases to 2. Hence, $p_1(\mathbf{r}, \mathbf{s}) = 5$. A similar reasoning applies to consumer 2 so that $p_2(\mathbf{r}, \mathbf{s}) = 5$. Consumer 3 is unmatched for reports below 4 and thus $p_3(\mathbf{r}, \mathbf{s}) = 4$.*[16]

Next, let us consider the first producer $j = 1$. In this case, if producer 1 reports a cost above 7, it would be unmatched by the Flip Algorithm. Hence, $w_1(\mathbf{r}, \mathbf{s}) = 7$. Producers 2 and 3 would be unmatched for reports above 9, which gives $w_2(\mathbf{r}, \mathbf{s}) = w_3(\mathbf{r}, \mathbf{s}) = 9$.

2.3.2 Matching Maximization

A second key property of the Flip Algorithm is that it yields a maximal number of value-creating matches, which we refer to as a *maximal matching*. As both consumers and producers are willing to report truthfully (Section 2.3.1), truth-telling is assumed in this subsection, *i.e.*, we assume that $\mathbf{r} = \mathbf{v}$ and $\mathbf{s} = \mathbf{c}$.

Definition 2.3.8 *Let $\mathcal{T}_{BR}$ be the set of all bilaterally rational matchings. A matching $T \in \mathcal{T}_{BR}$ is **maximal** when*

$$|T| = \max_{Q \in \mathcal{T}_{BR}} |Q|.$$

The set of all maximal matchings is given by $\mathcal{T}_M$.

The next example illustrates these concepts.

Example 2.3.9 *Consider the market of Example 2.3.1. Observe that producer 4 cannot engage in a value-creating match and that this market has a maximum of 3 value-creating matches. Therefore, $|T| = 3$ for any maximal matching T. Notice, however, that $A(\mathbf{r}, \mathbf{s}) = (\{1, 2, 3\}, \{1, 2, 3\})$ need not yield 3 value-creating matches. For example, matching $T_1 = \{(1, 3), (2, 2), (3, 1)\}$ or $T_2 = \{(1, 2), (2, 3), (3, 1)\}$ does, but $T_3 = \{(1, 1), (2, 2), (3, 3)\}$ does not. As a result, $T_1, T_2 \in \mathcal{T}_M$ and $T_3 \notin \mathcal{T}_M$.*

The next main result shows that the Flip Algorithm yields a maximal matching. In fact, it identifies the optimal maximal matching in the sense that it selects the maximal matching with the highest total utility.

Theorem 2.3.10 *The Flip Algorithm yields a matching in $\mathcal{T}_M$. Moreover, the Flip Algorithm maximizes the total utility over all matchings in $\mathcal{T}_M$.*

[16] Observe that the Flip Mechanism uses personalized prices, in contrast to the VCG mechanism that uses a uniform price for all consumers (see the proof of Proposition 2.4.2).

To explain the intuition behind the above findings, consider Figure 2.1 which graphically illustrates the market as described in Example 2.3.1. The blue step function represents the consumers' valuations, while the red step function represents the producers' costs. To begin, note that shifting the red curve rightward so that it is not above the blue curve is equivalent to the operation of the Flip Algorithm. The number of matched pairs is exactly the largest number of steps for which the red curve remains (weakly) below the blue curve and bilateral rationality is violated once a step of the red curve is above the blue curve. This visual crossing shows why the resulting matching is maximal.

Next, Figure 2.1 can also be used to explain monotonicity (Theorem 2.3.6). If a consumer increases its reported valuation, the corresponding blue step moves upward. Note that this can only enlarge (or leave unchanged) the region where the red curve lies below the blue curve, thereby (weakly) increasing the number of feasible value-creating matches and that consumer's likelihood of being matched. In case a consumer is already matched, the upward shift does not eliminate the existing overlap, so that it remains matched. The intuition for producers is analogous.

Finally, the Flip Algorithm maximizes total utility among all maximal matchings, because it connects the consumers with the highest valuations with the producers with the lowest costs.

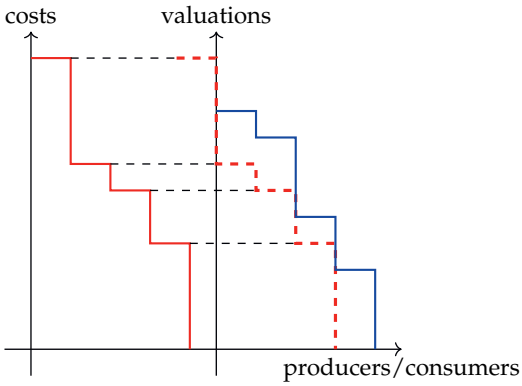

Figure 2.1: Illustration of the intuition underlying the Flip Algorithm using Example 2.3.1.

It is worth noting that there exists other algorithms to identify maximal matchings. Yet, the allocation rules resulting from such algorithms typically do not induce truth-telling.

Remark 2.3.11 *There exist so-called 'greedy algorithms' capable of identifying a maximal matching. Yet, allocation rules resulting from such algorithms are typically not monotonic and, therefore, not implementable. To illustrate, consider the following simple algorithm:*

1. *Order all consumers' reported valuations in a decreasing order.*
2. *Order all producers' reported costs in a decreasing order.*
3. *Starting with consumer* 1*, match consumer* i *to the first producer* j *for which it holds that* $r_i \geqslant s_j$*. Repeat this process, while bypassing any buyer that cannot be matched.*

Consider a market with 2 consumers and 2 producers. Table 2.6 presents their reported valuations and costs in an ordered fashion. The greedy algorithm matches consumer 1 to producer

r_1	r_2
10	5
s_1	s_2
8	6

Table 2.6: A market with 2 consumers and 2 producers.

1 and then stops. Suppose, however, that producer 2 would report a cost of 9 *instead. In that case, the producer will be matched, which violates monotonicity.*

2.4 The Price of Matching Maximization

Thus far, we have introduced an algorithm to identify the optimal maximal matching and shown implementability of the associated allocation rule. One may then wonder what it costs to actually implement it. To address this question, we compare the Flip Algorithm to the conventional TU-maximizing algorithm along two dimensions: the amount of deficit and the number of matches.

2.4.1 Total Utility Maximization: A Benchmark

The *TU-maximizing Algorithm* identifies a matching such that the total utility of all agents combined is maximal. It works as follows.

Recall that the reports of all consumers are ordered in decreasing order, *i.e.*, $r_1 \geqslant r_2 \geqslant \ldots \geqslant r_m$, and that the reports of all producers are ordered in increasing order, *i.e.*, $s_1 \leqslant s_2 \leqslant \ldots \leqslant s_n$. Let k be an index such that $r_l > s_l$ implies $l \leqslant k$ and $r_l < s_l$ implies $l > k$. Match k pairs $\{(i, i) \mid 1 \leqslant i \leqslant k\}$.

To see that this rule indeed maximizes total utility, note that $r_l \leqslant s_l$ for any $l > k$ so that matching k pairs generates (weakly) more utility than matching more than

k pairs. Moreover, $r_l \geqslant s_l$ for any $l \leqslant k$ so that matching up to k pairs (weakly) increases the total value-created. Note further that the TU-maximizing Algorithm yields a matching that is bilaterally rational. Indeed, since $r_1 \geqslant r_2 \geqslant \ldots \geqslant r_m$, $s_1 \leqslant s_2 \leqslant \ldots \leqslant s_n$ and the first k pairs are matched, it holds that $r_i \geqslant r_k \geqslant s_k \geqslant s_j$, for all $i \leqslant k$ and $j \leqslant k$.

Example 2.4.1 *To illustrate, consider the market as described in Example 2.3.1 above. In this example, $k = 2$ because the reported cost $s_3 = 7$ exceeds the reported valuation $r_3 = 5$ for the third pair. Therefore, $A = (\{1, 2\}, \{1, 2\})$ and the TU-maximizing Algorithm yields 2 pairs.*

It can be easily verified that a TU-maximizing allocation rule is monotonic and, therefore, implementable. In fact, as is well-known, one can use a VCG-mechanism to achieve an allocatively efficient outcome. Yet, this type of mechanism is typically not budget-balanced.[17] The next two results confirm that a VCG-mechanism indeed (almost) always creates a deficit, also in our setting.

Proposition 2.4.2 *If the policymaker adopts a VCG-mechanism, then $\sum_{i=1}^{m} p_i \leqslant \sum_{j=1}^{n} w_j$.*

Proposition 2.4.3 *The VCG-mechanism yields a deficit if, and only if, the allocation rule is unique.*

Taken together, this means that total utility maximization (almost) always requires external funding (*e.g.*, a subsidy).[18]

2.4.2 How (In)efficient is Matching Maximization?

Using the TU maximization results as reference, we now address the question of what it costs to maximize the number of value-creating matches. We consider two measures of efficiency: (i) the amount of deficit, and (ii) the number of matches. Let us discuss each in turn.

Amount of Deficit

How (in)efficient is matching maximization in monetary terms? The following theorem sheds some first light on this question.

Theorem 2.4.4 *Consider two implementable allocation rules A and A' where $K \supseteq K'$ and $L \supseteq L'$, for all $\mathbf{r} \in R$ and $\mathbf{s} \in S$. Denote the corresponding payment rules by p and p', respectively. Then, $p_i \leqslant p'_i$ for all consumers i and $w_j \geqslant w'_j$ for all producers j.*

[17] See, *e.g.*, Vickrey (1961), Myerson and Satterthwaite (1983) and Loertscher and Mezzetti (2019).

[18] As shown by McAfee (1992), one way to achieve budget-balancedness is by sacrificing one value-creating match.

The implication of this result is that the amount of deficit is (weakly) larger in case of matching maximization. This is because an implementable matching-maximizing allocation rule yields at least the same number of matches as the TU-maximizing allocation rule. Therefore, each consumer pays less, whereas each producer gets more. The following corollary summarizes this finding.

Corollary 2.4.5 *The deficit under an implementable matching-maximizing mechanism is at least as large as under a VCG-mechanism.*

It is worth emphasizing that this may hold even when the number of matches is the same under both mechanisms. This is illustrated by the next example.

Example 2.4.6 *Consider a market with 4 consumers and 4 producers. Table 2.7 contains the reported valuations and costs. Note that both the VCG Mechanism and the Flip Mechanism*

r_1	r_2	r_3	r_4
9	8	7	1
s_1	s_2	s_3	s_4
2	3	4	10

Table 2.7: A market with 4 consumers and 4 producers.

yield 3 matches. Under VCG, the deficit is 9 since each matched consumer pays 4 and each matched producer obtains 7. Under the Flip Mechanism, the deficit is 21 since each matched consumer pays 2 and each matched producer obtains 9. Hence, the deficit under the Flip Mechanism is larger than under the VCG Mechanism even though both give rise to 3 matches.

Knowing that matching maximization creates a (larger) deficit, one may wonder about its magnitude. How big can the deficit be? The next proposition provides an upper bound.

Proposition 2.4.7 *Suppose that $|T_F| = k$. The deficit is at most $k \cdot (r_1 - s_1)$.*

Number of Matches

An alternative way to assess the (in)efficiency of matching maximization is by analyzing the difference in the number of transactions. As the example in the introduction of this chapter illustrates, matching maximization may yield twice the number of transactions in comparison to TU maximization. As the next proposition shows, this is also the maximum ratio that can be attained.[19]

[19] This result also follows from known results in graph theory that the size of a maximal matching is at least half the size of a maximum matching.

Proposition 2.4.8 *The ratio of the number of matched pairs under matching maximization to the number of matched pairs under TU maximization is at most 2.*

It is worth noting that this finding extends to an incomplete bipartite graph. To see this, suppose that potential trading partners face some constraints. For example, a buyer and seller located far away from each other should perhaps not be matched since transportation costs would be too high. Imagine, then, an incomplete bipartite graph that indicates all feasible matches. To identify a maximal matching, one can proceed as follows. Assign flow 1 to each edge between consumer i's vertex and producer j's vertex when $r_i - s_j \geqslant 0$ and remove the edge when $r_i - s_j < 0$. A maximal matching on the resulting graph can then be found by applying Ford-Fulkerson's algorithm.

To identify a matching that maximizes total utility, one can proceed in a similar fashion. Assign flow $f_{ij} = r_i - s_j$ to each edge between consumer i's vertex and producer j's vertex when $r_i - s_j \geqslant 0$ and remove the edge when $r_i - s_j < 0$. One can then select the weighted maximum matching for the resulting graph. Now consider the resulting graph under matching maximization. To add one more match that maximizes the weight, one has to eliminate at most two matched pairs. Matching maximization will therefore not yield more than twice the number of matches under TU maximization.

2.5 Concluding Remarks

In the same year that Adam Smith launched his *Wealth of Nations* (1776), the prominent British scholar Jeremy Bentham published *A Fragment on Government* (1776). It is in this essay that Bentham formulated what became the fundamental maxim of utilitarianism, namely: "it is the greatest happiness of the greatest number that is the measure of right and wrong". Of course, when taken literally, the constructs 'greatest happiness' and 'greatest number' need not be reconcilable and pursuing one may well come at the cost of the other. While welfare economics has been predominantly concerned with the first by exploring traits and conditions that maximize the 'size of the pie', we focused on the second by asking how to maximize 'the number of pie bakers'.

Under the assumption that no market player can be forced to trade, we introduced an algorithm that identifies a matching with the maximum number of bilateral transactions. More specifically, we showed that this algorithm selects the optimal maximal matching in the sense that it creates the greatest happiness *conditional on* the greatest number. Importantly, the objective of matching maximization is implementable and we presented a mechanism that implements it. Doing so literally comes at a price, however, because, just like the VCG-mechanism in this setting, it does not satisfy the property of budget balance. A policymaker thus needs to find external funding if it seeks to maximize the number of matches.

Open questions remain regarding the implementation of the Flip Algorithm in more complex environments. Chapter 3 takes a step in this direction by examining a setting with unit-demand buyers and a monopolistic seller. Extending the analysis to multiple sellers, each capable of supplying multiple units, facing multi-unit demand remains open. Chapter 4 instead relaxes the assumption of a homogeneous good, considering markets with heterogeneous products. This setting is closely related to the assignment model of Demange et al. (1986), though Chapter 4 studies static allocation rather than the dynamic price-adjustment process analyzed there.

Appendix A: A Justification for Maximal Matchings

With its focus on maximizing the number of traded goods or services, matching maximization has an egalitarian flavor. The goal of this appendix is to make this more precise. In what follows, we consider and compare matchings' surplus distributions. It is shown that one obtains the most egalitarian outcome by applying the Flip Algorithm and have each matched pair divide its surplus equally.

To facilitate the analysis, we start with the following definition.

Definition A.1: *Consider a matching T. For each match $(i,j) \in T$, let $r_i - s_j = g_i + h_j$, where g_i and h_j are the surplus shares allocated to consumer i and producer j, respectively. Moreover, $g_i = 0$ for any unmatched consumer and $h_j = 0$ for any unmatched producer.*

Clearly, there are potentially many ways in which the surplus can be allocated. In what follows, let $\mathbf{d} = (g_1, \ldots, g_m, h_1, \ldots, h_n)$ be the surplus distribution vector. The *ordered* surplus distribution vector of $\mathbf{d}$ is $\hat{\mathbf{d}} = (d_1, \ldots, d_{m+n})$, where $d_i \leqslant d_j$ for all $1 \leqslant i \leqslant j \leqslant m+n$. Hence, the elements of $\hat{\mathbf{d}}$ are organized in a non-decreasing order.

To compare ordered surplus distributions, we use the following definition.

Definition A.2: *Consider two vectors $\mathbf{d}, \mathbf{e} \in \mathbb{R}^{m+n}$. The vector $\mathbf{d}$ is **lexicographically equal** to $\mathbf{e}$ when $\mathbf{d} = \mathbf{e}$. The vector $\mathbf{d}$ is **lexicographically larger** than $\mathbf{e}$ when $d_i > e_i$ for $i = \min\{j \in \{1, \ldots, m+n\} \mid d_j \neq e_j\}$. That is, $\mathbf{d}$ is lexicographically larger than $\mathbf{e}$ when the first coordinate that differs is higher under $\mathbf{d}$.*

Let us now establish the link between matching maximization and the egalitarian rule. Specifically, we show that the ordered surplus distribution that one obtains by maximizing the minimum utility is the same as the ordered surplus distribution that one obtains by applying the Flip Algorithm and allocating the resulting surpluses equally.

Proposition A.3 *Suppose that $r_i > s_j$ for all $(i,j) \in T_F$. The ordered surplus distribution that results from the Flip Algorithm $T_F(\mathbf{r}, \mathbf{s})$, with $g_i = h_j = \frac{r_i - s_j}{2}$ for any $(i,j) \in T_F$, is lexicographically larger or equal than the ordered surplus distribution of any other matching $T(\mathbf{r}, \mathbf{s})$.*

Appendix B: Proofs

Given the reported valuations of all consumers other than i and the reported costs of all producers, we write $a_i(r_i, \mathbf{r}_{-i}, \mathbf{s})$ as $a_i(r_i)$ and $a_i(t_i, \mathbf{r}_{-i}, \mathbf{s})$ as $a_i(t_i)$, for all $r_i, t_i \in \mathbb{R}_+$. In the same manner, we write $p_i(r_i, \mathbf{r}_{-i}, \mathbf{s})$ and $p_i(t_i, \mathbf{r}_{-i}, \mathbf{s})$ as $p_i(r_i)$ and $p_i(t_i)$, respectively for all $r_i, t_i \in \mathbb{R}_+$. Likewise, for the producers, we write $b_j(\mathbf{r}, s_j, \mathbf{s}_{-j})$ as $b_j(s_j)$ and $b_j(\mathbf{r}, u_j, \mathbf{s}_{-j})$ as $b_j(u_j)$, for all $s_j, u_j \in \mathbb{R}_+$. For the payment, denote $w_j(\mathbf{r}, s_j, \mathbf{s}_{-j})$ and $w_j(\mathbf{r}, u_j, \mathbf{s}_{-j})$ as $w_j(s_j)$ and $w_j(u_j)$.

Proof of Lemma 2.3.5

We first show that an implementable allocation rule is monotone. If the true valuation is r_i and the mechanism is DSIC, then it holds that:

$$r_i \cdot a_i(r_i) - p_i(r_i) \geqslant r_i \cdot a_i(t_i) - p_i(t_i) \iff p_i(t_i) - p_i(r_i) \geqslant r_i \cdot a_i(t_i) - r_i \cdot a_i(r_i). \tag{2.1}$$

Likewise, if the true valuation is t_i, then:

$$p_i(t_i) - p_i(r_i) \leqslant t_i \cdot a_i(t_i) - t_i \cdot a_i(r_i). \tag{2.2}$$

Combining (1) and (2), implies monotonicity for consumers. The proof for producers is analogous.

We now argue that an implementable allocation rule induces the payment rule given. This implies by the above proof that the allocation rule is monotonic. Using (1) and (2), observe that $p_i(r_i) = p_i(t_i)$ if $a_i(r_i) = a_i(t_i)$. So we need to find two prices: $p_i(t_i) = p_1$ if $a_i(t_i) = 0$ and $p_i(t_i) = p_2$ if $a_i(t_i) = 1$. Assume that $p_1 = 0$ (this is where the statement "up to an additive constant" comes from). Consider $\bar{t_i} = \inf\{t_i \mid a_i(t_i) = 1\}$. For all $t_i < \bar{t}_i$, we have $a_i(t_i) = 0$ and for all $t_i > \bar{t}_i$, we have $a_i(t_i) = 1$. Then by considering the left and right limit at $\bar{t}_i$, (1) and (2) imply the result. The argument for producers is analogous.

It remains to show that a monotone allocation rule together with the payment rule is indeed DSIC. Case 1: Consumer i receives a good by reporting their true valuation v_i. If $r_i > v_i$, then by monotonicity the consumer still receives a good and by the payment rule at the same price. If $r_i < v_i$, then either the consumer receives a good at the same price or the consumer does not receive a good, which induces a lower utility.

Case 2: Consumer i receives no good by reporting their true valuation v_i. If $r_i < v_i$, then by monotonicity the consumer still receives no good. If $r_i > v_i$, then either the consumer still receives no good or the consumer receives a good at a price higher than their valuation, which induces a lower utility. So it is a dominant strategy to report truthfully. The argument for producers is analogous. ■

Proof of Theorem 2.3.6

Recall that $|T_F| = \bar{k}$ and suppose that $a_i(r_i) = 1$. Fixing all other valuations and costs, we show that if consumer i reports $t_i \geqslant r_i$ instead, then $a_i(t_i) = 1$.

Since $a_i(r_i) = 1$, consumer i belongs to the first $\bar{k}$ matched consumers. If consumer i would report $t_i \geqslant r_i$ instead, then there are at least $\bar{k}$ matched pairs. To see this, consider a maximal matching $T \in \mathcal{T}_M$ in which consumer i is matched with producer j. Since all other valuations and costs remain the same, there still are $\bar{k} - 1$ value-creating matches when consumer i raises its reported valuation to $t_i \geqslant r_i$. As this consumer is matched when reporting r_i, it holds that $r_i \geqslant s_j$. Consequently, this consumer can be matched with the same producer when reporting t_i instead, because $t_i \geqslant r_i \geqslant s_j$. Therefore, the number of matches does not decrease when a consumer increases its reported valuation, all else equal.

As consumer i belongs to the first $\bar{k}$ matched pairs when reporting r_i, it also belongs to the first $\bar{k}$ matched pairs when reporting $t_i \geqslant r_i$. Therefore, $a_i(t_i) = 1$.

The argument for producers is analogous and thus we conclude that the allocation rule A_F is monotonic. ■

Proof of Theorem 2.3.10

Recall that we assume $\mathbf{r} = \mathbf{v}$ and $\mathbf{s} = \mathbf{c}$. We show that the Flip Algorithm induces a maximal matching and maximizes the total utility among all maximal matchings.

Let T be a maximal matching of size k. Let K and L denote the sets of consumers and producers matched under T, respectively. Write $K = \{i_1, \ldots, i_k\}$ such that $r_{i_t} \geqslant r_{i_{t+1}}$, and $L = \{j_1, \ldots, j_k\}$ such that $s_{j_t} \leqslant s_{j_{t+1}}$. Consider the matching $T' = \{(i_t, j_{k+1-t}) \mid 1 \leqslant t \leqslant k\}$.

Claim B.1 *T' is bilaterally rational.*

Proof. Consider consumer i_1 and producer j_k. Let j' denote the producer matched with consumer i_1 under T, and let i' denote the consumer matched with producer j_k under T.

If $i' = i_1$ (and hence $j_k = j'$), the claim follows immediately from the bilateral rationality of T, which implies $r_{i_1} \geqslant s_{j_k}$.

If $i' \neq i_1$, since i_1 has the highest valuation among all consumers in K, we have $r_{i_1} \geqslant r_{i'}$. Because (i', j_k) is a match in T, bilateral rationality implies $r_{i'} \geqslant s_{j_k}$. Hence, $r_{i_1} \geqslant s_{j_k}$. Moreover, since j_k has the highest cost among producers in L, we have $s_{j_k} \geqslant s_{j'}$. Combining this with $r_{i'} \geqslant s_{j_k}$ yields $r_{i'} \geqslant s_{j'}$.

We can therefore replace the pairs (i_1, j') and (i', j_k) with (i_1, j_k) and (i', j') without violating bilateral rationality. Remove i_1 and j_k from K and L, update the matching accordingly, and repeat this argument for the remaining agents. In the end, we obtain

a bilaterally rational matching $T' = \{(i_t, j_{k+1-t}) \mid 1 \leqslant t \leqslant k\}$. End proof of Claim 5.1.

From now on we assume that T has been transformed into T', as defined above. Order the (reported) valuations of all consumers in decreasing order, and the (reported) costs of all producers in increasing order. Formally, we have $r_1 \geqslant r_2 \geqslant \ldots \geqslant r_m$ and $s_1 \leqslant s_2 \leqslant \ldots \leqslant s_n$. Consider the matching $\hat{T} = \{(t, k+1-t) \mid 1 \leqslant t \leqslant k\}$, which is the outcome of the Flip Algorithm.

Claim B.2 *$\hat{T}$ is bilaterally rational and $TU(\hat{T}) \geqslant TU(T)$.*

Proof. We first establish the following auxiliary claim.

Claim B.2.1 *$r_t \geqslant r_{i_t} \geqslant s_{j_{k+1-t}} \geqslant s_{k+1-t}$, for all $1 \leqslant t \leqslant k$.*

Proof of Claim B.2.1 Fix a $t \in \{1, \ldots, k\}$. By Claim B.1, since $(i_t, j_{k+1-t}) \in T'$ and T' is bilaterally rational, we know that $r_{i_t} \geqslant s_{j_{k+1-t}}$. Moreover, since valuations are ordered in decreasing order, r_t is the t-th highest valuation among all consumers, while r_{i_t} is the t-th highest valuation among those matched under T. Hence, $r_t \geqslant r_{i_t}$ for all $1 \leqslant t \leqslant k$. Similarly, since costs are ordered in increasing order, s_{k+1-t} is the $(k+1-t)$-th lowest cost among all producers, while $s_{j_{k+1-t}}$ is the $(k+1-t)$-th lowest cost among matched producers, implying $s_{j_{k+1-t}} \geqslant s_{k+1-t}$. Hence, $r_t \geqslant r_{i_t} \geqslant s_{j_{k+1-t}} \geqslant s_{k+1-t}$. End proof of Claim B.2.1.

A. We now argue that $\hat{T}$ is bilaterally rational. It follows from Claim B.2.1 that $r_t \geqslant s_{k+1-t}$ for all $1 \leqslant t \leqslant k$. Hence, each match $(t, k+1-t)$ is value-creating. Therefore, $\hat{T}$ is bilaterally rational.

B. We show that $TU(\hat{T}) \geqslant TU(T)$. By Claim B.2.1, we have $r_{i_t} - s_{j_{k+1-t}} \leqslant r_t - s_{k+1-t}$, for all $1 \leqslant t \leqslant k$. Summing over $t = 1, \ldots, k$, we obtain $TU(\hat{T}) \geqslant TU(T)$. End proof of Claim B.2. ■

Proof of Proposition 2.4.2

Following the generic logic of VCG-mechanisms, consumers pay their externality. To determine the price for some matched consumer i, suppose that the TU-maximizing Algorithm yields k matches and is applied again without this consumer. If $r_{k+1} \geqslant s_k$, then the producer that was matched with consumer i before will be matched again. In this case, there are still k matches. If $r_{k+1} < s_k$, then the exclusion of consumer i implies the exclusion of some producer j. In this case, there are $k-1$ matches. Taken together, this means that any matched consumer pays a price $p_i = \max\{r_{k+1}, s_k\}$. By the same token, each matched producer obtains a price $w_j = \min\{s_{k+1}, r_k\}$.

Since the TU-maximizing Algorithm yields k matches, it holds that $s_k \leqslant r_k$ and $r_{k+1} < s_{k+1}$. Hence, $\max\{r_{k+1}, s_k\} \leqslant \min\{s_{k+1}, r_k\}$ and therefore:

$$\sum_{i=1}^{m} p_i - \sum_{j=1}^{n} w_j = k \cdot \max\{r_{k+1}, s_k\} - k \cdot \min\{s_{k+1}, r_k\} \leqslant 0.$$

Consequently, if the policymaker adopts a VCG-mechanism, then $\sum_{i=1}^{m} p_i \leqslant \sum_{j=1}^{n} w_j$. ■

Proof of Proposition 2.4.3

Suppose that the VCG-mechanism yields k matches. Let us first show that the allocation rule is unique when there is a deficit. To that end, assume that the allocation rule is not unique. We derive a contradiction. If the allocation rule is not unique, then there are at least two agents with rank k or $k+1$ that report the same value. Following the proof of Proposition 2.4.2, any matched consumer i pays $p_i = \max\{r_{k+1}, s_k\}$ and any matched producer j obtains $w_j = \min\{r_k, s_{k+1}\}$. Using the fact that $r_{k+1} \leqslant r_k$, $r_{k+1} < s_{k+1}$, $s_k \leqslant s_{k+1}$, and $s_k \leqslant r_k$, this effectively leaves two possibilities.

1. If $s_k \leqslant r_k = r_{k+1} < s_{k+1}$, then $p_i = \max\{r_{k+1}, s_k\} = r_{k+1}$ and $w_j = \min\{r_k, s_{k+1}\} = r_k$. Hence, $p_i = w_j$ for all $i \in K$ and $j \in L$. In this case, therefore, there is no deficit.

2. If $r_{k+1} < s_k = s_{k+1} \leqslant r_k$, then $p_i = \max\{r_{k+1}, s_k\} = s_k$ and $w_j = \min\{r_k, s_{k+1}\} = s_{k+1}$. Hence, $p_i = w_j$ for all $i \in K$ and $j \in L$. In this case, therefore, there is no deficit.

We conclude that if there is a deficit, then the allocation rule is unique.

Let us now show that there is a deficit when there is a unique allocation rule. Suppose, by contradiction, that there is a unique allocation rule, but no deficit. By the preceding analysis, all matched consumers pay $\max\{r_{k+1}, s_k\}$ and all matched producers obtain $\min\{r_k, s_{k+1}\}$. If $\max\{r_{k+1}, s_k\} = r_{k+1}$, then there is a deficit since $r_k > r_{k+1}$ and $s_{k+1} > r_{k+1}$. If $\max\{r_{k+1}, s_k\} = s_k$, then there is a deficit since $s_k < r_k$ by uniqueness and $s_k < s_{k+1}$. We conclude that if the allocation rule is unique and $r_k > s_k$, then there is a deficit. ■

Proof of Theorem 2.4.4

As A and A' are implementable, they are monotonic (Myerson's Lemma). Consider some consumer i with 'switch point' r_i under $A(\mathbf{r}, \mathbf{s})$ and 'switch point' t_i under $A'(\mathbf{r}, \mathbf{s})$. Let us now view the allocation of consumer i with valuation t_i under $A(\mathbf{r}, \mathbf{s})$. Since $K \supseteq K'$ for all $\mathbf{r} \in R$, $\mathbf{s} \in S$, it holds that $a_i(t_i) \geqslant a'_i(t_i)$ and, therefore, $t_i \geqslant r_i$. Since A is implementable by payment rule p, each consumer pays its 'switch point' where it goes from 'not being matched' to 'being matched'. Similarly for A'. Therefore, $p'_i = t_i \geqslant r_i = p_i$.

Now consider some producer j with 'switch point' s_j under $A(\mathbf{r}, \mathbf{s})$ and 'switch point' u_j under $A'(\mathbf{r}, \mathbf{s})$. Let us now view the allocation of producer j with cost u_j under $A(\mathbf{r}, \mathbf{s})$. Since $L \supseteq L'$ for all $\mathbf{r} \in R$, $\mathbf{s} \in S$, it holds that $b_j(u_j) \geqslant b'_j(u_j)$ and, therefore, $u_j \leqslant s_j$. Since A is implementable by payment rule p, each producer obtains its 'switch point' where it goes from 'not being matched' to 'being matched'. Similarly for A'. Therefore, $w'_j = u_j \leqslant s_j = w_j$. ■

Proof of Proposition 2.4.7

The lowest possible price for a matched consumer is s_1, whereas the highest possible price for a matched producer is r_1. Since there are k matches, the greatest possible deficit is $k \cdot (r_1 - s_1)$. ■

Proof of Proposition 2.4.8

To begin, note that the example in the introduction of this chapter shows that the ratio can be 2. Let us now show that it cannot be more than 2. To that end, suppose that the TU-maximizing Algorithm yields k matches, which implies $r_{k+m} < s_{k+m}$ for all $m \geqslant 1$. Suppose further that the Flip Algorithm yields $2k + m$ matches, where $m \geqslant 1$. We derive a contradiction. If the Flip Algorithm leads to $2k + m$ matches, then consumer 1 is matched with producer $2k + m$ and consumer 2 is matched with producer $2k + m - 1$. Following this logic, consumer $k + 1$ is matched with producer $k + m$, which implies $r_{k+1} \geqslant s_{k+m} \geqslant s_{k+1}$. This, however, contradicts the fact that $r_{k+1} < s_{k+1}$. ■

Proof of Proposition A.3

To begin, note that the ordered surplus distribution of any bilaterally rational matching is lexicographically larger than the ordered surplus distribution of a matching that is not bilaterally rational. In case of the latter, there is at least one agent that obtains a strictly negative surplus, whereas all agents receive a non-negative surplus when the matching is bilaterally rational. Hence, since the Flip Algorithm yields a bilaterally rational matching, the resulting ordered surplus distribution is lexicographically larger than the ordered surplus distribution of any matching that is not bilaterally rational.

Next, consider a bilateral matching T, where $|T| < |T_F|$. Since each match has a strictly positive surplus and unmatched agents receive zero surplus, it follows immediately that the ordered surplus distribution that results from the Flip Algorithm is lexicographically larger than the ordered surplus distribution of any matching with strictly fewer matches.

Finally, consider a matching T such that $|T| = |T_F| = k$ but $T \neq T_F$. Let K and L denote the sets of consumers and producers matched under T, respectively. Write $K = \{i_1, \ldots, i_k\}$ such that $r_{i_t} \geqslant r_{i_{t+1}}$, and $L = \{j_1, \ldots, j_k\}$ such that $s_{j_t} \leqslant s_{j_{t+1}}$.

We first define when two matches cross. Consider two distinct matches (i, j) and (i', j') in T. We say that (i, j) and (i', j') cross if $(i - i') \cdot (j - j') < 0$, that is, if they

are not strictly ordered in the same direction. Otherwise, they are uncrossed. Let $u(T)$ denote the number of uncrossed pairs of matches in T, that is,

$$u(T) = \big|\{\{(i,j),(i',j')\} \subseteq T \mid (i-i')\cdot(j-j') > 0\}\big|.$$

Consider the matching $S = \{(i_t, j_{k+1-t}) \mid 1 \leqslant t \leqslant k\}$. Observe that $u(T) = 0$ if and only if $T = S$.[20]

If $u(T) > 0$, consider an uncrossed pair of matches (i,j) and (i',j') in T. Without loss of generality, assume that $i > i'$. Since they are uncrossed, we know that $j > j'$. Define

$$T' = \big(T \backslash \{(i,j),(i',j')\}\big) \cup \{(i,j'),(i',j)\}.$$

We prove the following three claims.

Claim A.3.1 *T' is bilaterally rational.*

Proof of Claim A.3.1. Since (i,j) and (i',j') are in T and T is bilaterally rational, we know that $r_i \geqslant s_j$ and $r_{i'} \geqslant s_{j'}$. Furthermore, because $i > i'$ and $j > j'$, we also know that $r_{i'} \geqslant r_i$ and $s_{j'} \leqslant s_j$. Combining these inequalities, we obtain $r_{i'} \geqslant s_j$ and $r_i \geqslant s_{j'}$. Therefore, the matches (i',j) and (i,j') are value-creating. All other matches in T' coincide with those in T, and hence are also value-creating. Consequently, T' is bilaterally rational. End proof of Claim A.3.1.

Claim A.3.2 *The ordered surplus distribution of T' is lexicographically larger or equal to the ordered surplus distribution of T.*

Proof of Claim A.3.2. Consider the two matches (i,j) and (i',j') in T, and the corresponding matches (i,j') and (i',j) in T'. Since $i > i'$ and valuations are ordered decreasingly, we have $r_i \leqslant r_{i'}$. Similarly, since $j > j'$ and costs are ordered increasingly, we have $s_j \geqslant s_{j'}$. Hence, $r_i - s_j \leqslant r_{i'} - s_{j'}$. Moreover, since $s_{j'} \leqslant s_j$, it follows that $r_i - s_{j'} \geqslant r_i - s_j$, and since $r_{i'} \geqslant r_i$, we have $r_{i'} - s_j \geqslant r_i - s_j$. Thus, both new surpluses are (weakly) larger than the smallest surplus among the two original matches. If both are equalities, then $r_{i'} = r_i$ and $s_j = s_{j'}$, so the two surpluses are unchanged. Since all other matches remain unchanged, it follows that, after ordering surpluses increasingly, the smallest element (weakly) increases while all others remain the same. Hence, the ordered surplus distribution of T' is lexicographically larger or equal than that of T. End proof of Claim A.3.2

[20] If $T = S$, it is immediate that $u(T) = 0$. If $u(T) = 0$, then every pair of matches in T is crossed, that is, for any two matches (i,j) and (i',j') in T, we have $(i-i')\cdot(j-j') < 0$. This implies that matches are ordered in opposite directions. For all $i > i'$, we must have $j < j'$. Hence, larger-indexed consumers are always matched with smaller-indexed producers. Since T is a matching of size k, this ordering uniquely determines the assignment, and it follows that T must coincide with the matching $S = \{(i_t, j_{k+1-t}) \mid 1 \leqslant t \leqslant k\}$.

Claim A.3.3 $u(T') < u(T)$.

Proof of Claim A.3.3. We first consider a match $(i'', j'') \in T \backslash \{(i, j), (i', j')\}$ and check the number of crossed pair involved (i'', j''). Consider the following cases:

Case 1: $i'' < i'$ and $j'' < j'$. In this case, (i'', j'') crosses neither (i', j') nor (i, j) before the swap. After the swap, it also crosses neither (i', j) nor (i, j'). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

Case 2: $i'' < i'$ and $j' < j'' < j$. In this case, (i'', j'') crosses (i', j') and does not cross (i, j) before the swap. After the swap, it crosses (i, j') and does not cross (i', j). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

Case 3: $i'' < i'$ and $j'' > j$. In this case, (i'', j'') crosses both (i', j') and (i, j) before the swap. After the swap, it also crosses both (i', j) and (i, j'). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

Case 4: $i' < i'' < i$ and $j'' < j'$. In this case, (i'', j'') crosses (i', j') and does not cross (i, j) before the swap. After the swap, it crosses (i', j) and does not cross (i, j'). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

Case 5: $i' < i'' < i$ and $j' < j'' < j$. In this case, (i'', j'') crosses neither (i', j') nor (i, j) before the swap. After the swap, it crosses both (i', j) and (i, j'). Hence, the number of crossed pairs involving (i'', j'') strictly increases.

Case 6: $i' < i'' < i$ and $j'' > j$. In this case, (i'', j'') crosses (i, j) and does not cross (i', j') before the swap. After the swap, it crosses (i, j') and does not cross (i', j). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

Case 7: $i'' > i$ and $j'' < j'$. In this case, (i'', j'') crosses both (i', j') and (i, j) before the swap. After the swap, it also crosses both (i', j) and (i, j'). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

Case 8: $i'' > i$ and $j' < j'' < j$. In this case, (i'', j'') crosses (i, j) and does not cross (i', j') before the swap. After the swap, it crosses (i', j) and does not cross (i, j'). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

Case 9: $i'' > i$ and $j'' > j$. In this case, (i'', j'') crosses neither (i', j') nor (i, j) before the swap. After the swap, it also crosses neither (i', j) nor (i, j'). Hence, the number of crossed pairs involving (i'', j'') remains unchanged.

In all cases, the number of uncrossed pairs involving (i'', j'') does not increase. Combined with the fact that the pair $\{(i, j), (i', j')\}$ becomes crossed after the swap, we obtain that $u(T') < u(T)$. End proof of Claim A.3.3.

We now use the preceding claims to establish the following result.

Claim A.3.4 *The ordered surplus distribution of S is lexicographically larger or equal to that of T.*

Proof of Claim A.3.4. As long as $u(T) > 0$, we iteratively apply the transformation described above. By Claim A.3.2, each application (weakly) increases the ordered surplus distribution in the lexicographic order. By Claim A.3.3, each transformation strictly decreases the number of uncrossed pairs. Hence, the process terminates at a matching T'' with $u(T'') = 0$. By the observation above, $T'' = S$. It follows that the ordered surplus distribution of S is lexicographically larger or equal to that of T. End proof of Claim A.3.4.

It remains to compare the ordered surpluses of S and $\hat{S} = \{(t, k+1-t) \mid 1 \leqslant t \leqslant k\}$, which is the outcome of the Flip Algorithm. Define the surpluses of S by $x_t = r_{i_t} - s_{j_{k+1-t}}$ and of $\hat{S}$ by $y_t = r_t - s_{k+1-t}$, for all $1 \leqslant t \leqslant k$. Let $\sigma : \{1, \ldots, k\} \mapsto \{1, \ldots k\}$ be a permutation such that $x_{\sigma(1)} \leqslant x_{\sigma(2)} \leqslant \ldots \leqslant x_{\sigma(k)}$. By Claim B.2.1, we have that $y_{\sigma(t)} \geqslant x_{\sigma(t)}$ for all $1 \leqslant t \leqslant k$. Consider the smallest surplus $y_{\sigma(\ell)}$, i.e., $y_{\sigma(\ell)} \leqslant y_{\sigma(t)}$ for all $1 \leqslant t \leqslant k$. Define $\hat{\mathbf{y}}$ by

$$\hat{y}_t = \begin{cases} y_{\sigma(\ell)} & \text{if } t = 1, \\ y_{\sigma(t-1)} & \text{if } 2 \leqslant t \leqslant \ell, \\ y_{\sigma(t)} & \text{if } \ell + 1 \leqslant t \leqslant k. \end{cases}$$

We now prove that $\hat{y}_t \geqslant x_{\sigma(t)}$ for all $1 \leqslant t \leqslant k$. First, consider $\hat{y}_1$, which is the smallest surplus of $\hat{S}$. Then $\hat{y}_1 = y_{\sigma(\ell)} \geqslant x_{\sigma(\ell)} \geqslant x_{\sigma(1)}$. Second, consider $\hat{y}_t$ for all $2 \leqslant t \leqslant \ell$. Then $\hat{y}_t \geqslant \hat{y}_1 = y_{\sigma(\ell)} \geqslant x_{\sigma(\ell)} \geqslant x_{\sigma(t)}$. Finally, consider $\hat{y}_t$ for all $\ell + 1 \leqslant t \leqslant k$. We know that $\hat{y}_t = y_{\sigma(t)} \geqslant x_{\sigma(t)}$.

Observe that the smallest surplus of S, $x_{\sigma(1)}$, is smaller or equal to that of $\hat{S}$, $\hat{y}_1 = y_{\sigma(\ell)}$, so we can delete $x_{\sigma(1)}$ and $y_{\sigma(\ell)}$, and apply the above argument to the remaining surpluses. In the end, we conclude that the resulting surplus vector of $\hat{S}$ is lexicographically larger or equal to that of S. ■

3

Monopolistic Trade Maximization

This chapter is based on joint work with Iwan Bos, Marc Schröder and Dries Vermeulen (Maastricht University).

Earlier versions of this work were presented at the SING20 conference, Joint ORBEL–NGB conference, the Dutch Game Theory Symposium, the MLSE at Maastricht University and the Microeconomics Seminar at University of Lausanne.

3.1 Introduction

The question of how to optimally allocate scarce resources lies at the heart of economic science. In assessing whether a particular allocation can be considered 'optimal' it needs to be clear what the desired outcome is. In the language of mechanism design, what *social goal* should be pursued?[1] Perhaps the most commonly known social goal is captured by the so-called utilitarian rule. This rule states that the preferred alternatives are those that maximize the aggregate utility of all agents involved. Yet, recent work in the field of mechanism design puts an increasing emphasis on other potentially important objectives, most notably related to issues of equity and fairness.[2] This raises concerns about how much allocative surplus needs to be sacrificed when taking account of (re-)distribution considerations.

In this chapter, we also consider an alternative social goal, namely the maximization of trade volume. A prime rationale for this goal is that a policymaker may be concerned about the availability of essential goods or services such as electricity, housing, medical supplies and transportation services. Note that there would be little need to deviate from the utilitarian rule when general, unconstrained transfers are possible. In many markets, however, transactions are inherently bilateral (*e.g.*, taxi services), and transfers across agents may be difficult or impossible to implement. For instance, there is no cross-subsidization among taxi drivers.

Specifically, we study two settings in which a monopolist serves (*i*) a set of consumers with unit demand, and (*ii*) one consumer with multi-unit demand (bilateral monopoly).[3] Both the costs of production and the valuations for the product are private information. Within this context, we examine whether and under what conditions there are implementable, value-creating allocation rules that yield more transactions than would take place under the utilitarian rule. In turn, this sheds light on the potential (efficiency) cost of aiming for trade rather than total utility maximization.

Our main findings are as follows. First, we show that the utilitarian rule is undominated in terms of the number of transactions when one does not impose any assumption on the cost function (other than that marginal and total production costs are positive). That is to say, there exists no implementable, value-creating allocation rule that consistently yields more transactions than are obtained with total utility maximization. Second, it is possible to go beyond the total utility maximizing outcome when marginal cost is non-decreasing in output and we introduce an incentive compatible rule, coined Average Allocation Rule (AAR), that achieves this.[4] The AAR not only

[1] See, *e.g.*, Maskin (2008).

[2] See, *e.g.*, Condorelli (2013), Dworczak et al. (2021), Groh and Reuter (2023), Akbarpour et al. (2024) and Finster et al. (2024).

[3] We have studied the case of multiple producers, each of which can supply a single item in Chapter 2. Chapter 4 then takes the same question to markets with heterogeneous products, where objects are no longer interchangeable.

[4] In the ensuing analysis, it will become clear why we have labeled this rule *Average Allocation* Rule. Note further that convex cost functions are widely used in economics, *e.g.*, to describe situations in which a firm faces capacity constraints. Correa et al. (2014) for example, describe how an electricity company initially

dominates the utilitarian rule, but is also undominated by any other implementable, value-creating allocation rule in that it yields a maximal number of matches. Yet, this rule is not bilaterally value-creating, meaning that not every individual transaction generates a surplus (*i.e.*, there is cross-subsidization across transactions). In fact, our third main finding is that there exists no implementable, bilaterally value-creating allocation rule that always yields more matches than result under total utility maximization. Fourth and finally, we establish the equivalence between total utility maximization and trade maximization in the case of bilateral monopoly with multi-unit demand. Taken together, these findings show that the utilitarian rule performs quite well in terms of trade volume when there is a single supplier.

In the next section, the model is introduced. Section 3.3 considers total utility maximization without any particular assumption on the cost function. Section 3.4 provides a description and analysis of Average Allocation Rules. Section 3.5 is devoted to bilateral value-creation, whereas Section 3.6 considers the case of bilateral monopoly with multi-unit demand. Section 3.7 concludes. All proofs are relegated to the Appendix.

3.2 Model

There is a single supplier that can offer an amount of $m \geqslant 0$ units of an indivisible good (items). The total cost of producing k items, $k = 0, 1, \ldots, m$, is given by $c(k)$ and the marginal cost of producing the kth item is denoted by $c_k = c(k) - c(k-1)$. It is assumed that $c(0) = 0$ and that $c_k \geqslant 0$ so that total cost $c(k)$ is non-negative and non-decreasing in k. Demand comes from a set of consumers, $M = \{1, \ldots, m\}$. Each consumer $i \in M$ obtains at most one item, which it values at $v_i \geqslant 0$. The cost of production as well as the valuations are private information.

In what follows, we focus on transactions that are value-creating. To make this precise, let us define the notion of an allocation rule. An allocation rule is a function that assigns to each vector of valuations, $v = (v_1, \ldots, v_m) \in \mathbb{R}_+^m$, and each vector of marginal costs, $c = (c_1, \ldots, c_m) \in \mathbb{R}_+^m$, a set of consumers $A(v, c) \subseteq M$ that obtain an item:

$$A \colon \mathbb{R}_+^m \times \mathbb{R}_+^m \mapsto 2^M.$$

The $k = |A(v, c)|$ items are then produced at a total cost of $c(k) = \sum_{i=1}^{k} c_i$.

Definition 3.2.1 *An allocation rule A is **value-creating** if, for all (v, c),*

$$\sum_{i \in A(v,c)} v_i \geqslant c(|A(v, c)|).$$

utilizes its most efficient generators and activate less efficient ones only in case of excess demand.

Value-creation thus ensures that the allocation is economically viable in the sense that the combined valuation of consumers involved in a transaction weakly exceeds the associated production costs.

Next, we introduce the related concept of bilateral value-creation. Toward that end, let the set $T \subseteq M \times M$ be a matching that indicates which item the producer gives to which consumer. Specifically, a match $(i,k) \in T$ means that consumer i gets the kth item. Since the producer cannot offer the kth item without having supplied the $(k-1)^{th}$ item, $(i,k) \in T$ implies that there is a $j \in M$ with $(j,k-1) \in T$.

Definition 3.2.2 *A matching T is **bilaterally value-creating** when $v_i \geqslant c_k$ for each match $(i,k) \in T$.*

Definition 3.2.3 *An allocation rule A is **bilaterally value-creating** if, for all (v,c), there exists a bilaterally value-creating matching T such that:*

$$A(v,c) = \{i \in M \mid (i,k) \in T \text{ for some } k \in M\}.$$

Note that bilateral value-creation is a stronger condition than value-creation as it requires surplus generation at the level of each individual transaction rather than just at the aggregate level.

A natural question is now how many items are produced and how these items are allocated among the consumers. To address this issue, we assume there is a social planner that employs a direct mechanism to achieve its objective.[5] A direct mechanism is a triplet $\mathcal{M} = (A,p,w)$, where $A\colon \mathbb{R}_+^m \times \mathbb{R}_+^m \mapsto 2^M$ is the allocation rule, $p\colon \mathbb{R}_+^m \times \mathbb{R}_+^m \mapsto \mathbb{R}^m$ is the payment rule for the consumers and $w\colon \mathbb{R}_+^m \times \mathbb{R}_+^m \mapsto \mathbb{R}$ is the reward rule for the producer. As an input, the social planner receives a (non-negative) vector of reported valuations, $r = (r_1,\ldots,r_m) \in \mathbb{R}_+^m$, from the consumers and a (non-negative) vector of reported marginal costs, $s = (s_1,\ldots,s_m) \in \mathbb{R}_+^m$, from the producer. Given r and s, the set $A(r,s) \subseteq M$ is then the collection of consumers that obtain an item. For each consumer $i \in M$, the outcome of the market interaction is thus given by:

$$a_i(r,s) = \begin{cases} 1 \text{ if } i \in A(r,s) \\ 0 \text{ if } i \notin A(r,s). \end{cases}$$

In the ensuing analysis, we consider allocation rules that are implementable. For this reason, we now introduce some key concepts that are used to ensure incentive compatibility.

Definition 3.2.4 *A mechanism is **dominant strategy incentive compatible (DSIC)** if,*

[5] The use of a direct mechanism is justified by the so-called revelation principle (Gibbard (1973)), which states that when there is a mechanism that yields a particular outcome, there exists a direct incentive-compatible mechanism that delivers the same outcome.

[1] *for all* $i \in M$*, for all* $r_i, r_i' \in \mathbb{R}_+$*, for all* $r_{-i} \in \mathbb{R}_+^{m-1}$ *and* $s \in \mathbb{R}_+^m$*, it holds that*

$$r_i \cdot a_i(r_i, r_{-i}, s) - p_i(r_i, r_{-i}, s) \geqslant r_i \cdot a_i(r_i', r_{-i}, s) - p_i(r_i', r_{-i}, s), \quad \textit{and}$$

[2] *for the producer, for all* $s, s', r \in \mathbb{R}_+^m$*, it holds that*

$$w(r, s) - s(k) \geqslant w(r, s') - s(k'),$$

where $k = |A(r, s)|$ *and* $k' = |A(r, s')|$.

A mechanism is DSIC when the net utility of reporting the actual valuations or costs is weakly higher than the net utility of reporting anything else so that reporting truthfully is a dominant strategy for the consumers as well as the producer.

Definition 3.2.5 *An allocation rule* A *is* ***implementable*** *if there is a payment rule* (p, w) *such that* (A, p, w) *is DSIC.*

Within a multi-dimensional setting, Bikhchandani et al. (2003) showed that an allocation rule is implementable precisely when it is monotonic. A similar result regarding the equivalence of implementability and monotonicity was obtained by Lavi et al., 2003, although their proof takes a different approach. Since we use this equivalence extensively in this chapter, we now present the definition of monotonicity and the equivalence result within our setting.

Definition 3.2.6 *An allocation rule* A *is* ***monotonic*** *if*

[1] *for all* i*, and for all* $r, r', s \in \mathbb{R}_+^m$ *with* $r_{-i} = r'_{-i}$*,* $i \in A(r, s)$*, and* $i \notin A(r', s)$*, it holds that* $r_i > r_i'$*, and*

[2] *for all* $r, s, s' \in \mathbb{R}_+^m$*, it holds that*

$$\sum_{i=k'+1}^{k} s_i \leqslant \sum_{i=k'+1}^{k} s_i',$$

where $k = |A(r, s)|$ *and* $k' = |A(r, s')|$*, and* $k > k'$.

When an allocation rule is monotonic, a consumer will not receive less when reporting a higher valuation. Moreover, a change in the reported marginal cost of the kth item does not affect the production decision about any other item.

Lemma 3.2.7 *An allocation rule* A *is implementable if, and only if,* A *is monotonic.*

The consumer-side characterization follows from Myerson, 1981. The producer-side characterization follows from the weak monotonicity characterization of Bikhchandani et al. (2003), tailored to our setting.

In what follows, we use the above model to examine the relation between the maximization of total utility and the maximization of the number of transactions. As to the latter, we exclusively consider allocation rules that are monotonic and value-creating.

Definition 3.2.8 *Let $\mathcal{A}$ be the collection of all monotonic and value-creating allocation rules. An allocation rule $A \in \mathcal{A}$ is **maximal** if there does not exist an allocation rule $B \in \mathcal{A}$ such that, for all reports (r, s), it holds that:*

$$|A(r, s)| \leqslant |B(r, s)|$$

and the inequality is strict for at least one report (r, s).

The maximization of total utility is elaborated on in the next section.

3.3 Total Utility Maximization as Upper Bound

In this section, we start by exploring how the maximal number of matches under total utility maximization relates to the maximal number of value-creating matches. Toward that end, we first specify the concepts of *total utility* (TU) and *TU-maximizer* within our setting.

Definition 3.3.1 *For any given allocation $B \subseteq M$, total utility is defined as the difference between the combined valuations of all matched consumers and the cost of production:*

$$TU(B)(v, c) = \sum_{i \in B} v_i - c(|B|).$$

Definition 3.3.2 *An allocation rule A is a TU-maximizer if, for every allocation $B \subseteq M$ and every (v, c), it holds that:*

$$TU(A(v, c))(v, c) \geqslant TU(B)(v, c).$$

Note that there may be multiple TU-maximizers that can differ in terms of cardinality. In what follows, our focus is on TU-maximizers with the largest number of value-creating matches.

The next theorem establishes that total utility maximization is undominated regarding the number of value-creating matches. That is to say, any monotonic and value-creating allocation rule that yields at least the number of matches under TU maximization also yields at most the number of matches under TU maximization.

Theorem 3.3.3 *Let A be any monotonic and value-creating allocation rule and suppose that A^{TU} is a TU-maximizer with maximal cardinality. If $|A(r,s)| \geqslant |A^{TU}(r,s)|$ for all (r,s), then $|A(r,s)| = |A^{TU}(r,s)|$ for all (r,s).*

The proof exploits the absence of restrictions on the marginal cost function. In particular, if strictly more consumers are selected under A than under A^{TU}, the marginal cost of the first item can be increased while the marginal costs of subsequent items can be decreased in such a way that the additional matches selected by A create negative value. This forces A to drop some of these matches. However, this violates monotonicity, since the rule reduces the number of matches even though the relevant costs have decreased.

3.4 Average Allocation Rules

In the previous section, we established that the utilitarian rule is undominated in terms of the number of value-creating matches in that any monotonic and value-creating allocation rule that yields at least as many matches also yields at most this number of matches. This result was derived without hardly any restrictions on the production function. In particular, we allowed for reductions in marginal cost over ranges of output.

In this section, we suppose that the production technology is such that the producer's marginal cost is weakly increasing in output and assume this to be common knowledge. That is, both the actual and the reported marginal costs are non-decreasing (*i.e.*, $0 \leqslant c_1 \leqslant \ldots \leqslant c_m$ and $0 \leqslant s_1 \leqslant \ldots \leqslant s_m$). Under this assumption, we introduce a type of rule, coined *Average Allocation Rule* (AAR) that (*i*) dominates the total utility maximizing outcome in terms of the number of value-creating matches, and (*ii*) is itself undominated in the sense that it is maximal (Definition 3.2.8).

Let us begin by specifying an Average Allocation Rule.

Definition 3.4.1 *An allocation rule A is an Average Allocation Rule (AAR) if the following holds. For every (r,s), let $\sigma : M \mapsto M$ be such that $r_{\sigma(1)} \geqslant \ldots \geqslant r_{\sigma(m)}$. Then, for $k = |A(r,s)|$, it holds that*

1. *$A(r,s) = \{\sigma(1), \ldots, \sigma(k)\}$, and*
2. *$k = \max\left\{\ell \mid \frac{1}{\ell} \cdot \sum_{i=1}^{\ell} r_{\sigma(i)} \geqslant s_\ell\right\}$.*

The following example illustrates the workings of an AAR.[6]

$r_{\sigma(1)}$	$r_{\sigma(2)}$	$r_{\sigma(3)}$
8	4	3
s_1	s_2	s_3
1	5	7

Table 3.1: A market with three consumers and one producer.

Example 3.4.2 *Consider a market with three consumers and one producer. Table 3.1 depicts the reported valuations and marginal costs in an ordered fashion.*

According to an Average Allocation Rule, $A(r,s) = \{\sigma(1), \sigma(2)\}$ *and the producer produces 2 units since* $\frac{r_{\sigma(1)}+r_{\sigma(2)}}{2} = 6 > 5 = s_2$ *and* $\frac{r_{\sigma(1)}+r_{\sigma(2)}+r_{\sigma(3)}}{3} = 5 < 7 = s_3$.

The next result establishes that average allocation rules are value-creating (Definition 3.2.1) and implementable (Lemma 3.2.7).

Proposition 3.4.3 *Average allocation rules are value-creating and monotonic.*

Our second main finding is that average allocation rules dominate TU-maximizers in terms of the number of value-creating matches.

Theorem 3.4.4 *Suppose that* A^{TU} *is a TU-maximizer with maximal cardinality. If* A *is an Average Allocation Rule, then* $|A(r,s)| \geqslant |A^{TU}(r,s)|$ *for all* (r,s).

Remark 3.4.5 *Theorem 3.4.4 is a result on the number of value-creating matches and does not establish set inclusion between the allocations induced by an AAR and a TU-maximizer in general. In particular, when ties are present, the induced allocations need not be nested. For example, consider a market with two consumers and one producer. Table 3.2 illustrates the reported valuations and marginal costs.*

r_1	r_2
2	2
s_1	s_2
1	5

Table 3.2: A market with two consumers and one producer.

A TU-maximizer may allocate one unit to consumer 1, whereas an AAR may allocate one unit to consumer 2.

[6]Note that we do not refer to *the* AAR here since the possibility of tie-breaking results in different average allocation rules.

The finding of Theorem 3.4.4 is visualized in Figure 3.1. The blue curve gives the number of matches obtained under the utilitarian rule, whereas the green curve represents the number of matches resulting from an AAR. As can be observed, the green curve lies weakly above the blue curve for any realization of (r, s). By comparison, there may be some other allocation rule A, such as the one depicted by the red curve, that in some cases yields more and in some cases yields less value-creating matches.[7] Unlike an AAR, therefore, A does not dominate the TU-maximizer.

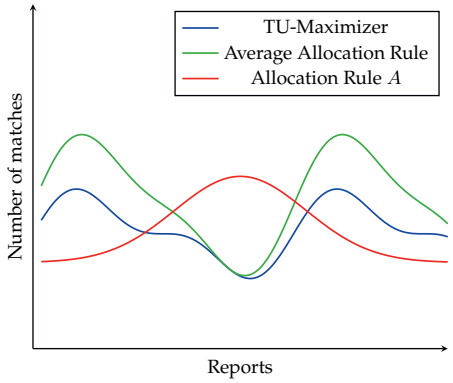

Figure 3.1: Theorem 3.4.4 illustrated.

The reason that an AAR may create strictly more matches than a TU-maximizer is that it cross-subsidizes across transactions. Indeed, it effectively takes some surplus from value-creating matches to facilitate additional ones that do not create value. Note that this is feasible only when the marginal cost curve is not too steep at the total utility maximizing outcome.[8] For instance, in Example 3.4.2 above, the AAR yields two matches, whereas total utility maximization is obtained with one transaction.

Finally, we show that average allocation rules are themselves undominated by any other monotonic and value-creating allocation rule in that they yield a maximum number of transactions. In other words, using the terminology of Definition 3.2.8 above, an AAR is maximal.

Theorem 3.4.6 *Let A be an Average Allocation Rule and B be a monotonic and value-creating allocation rule. If $|A(r,s)| \leqslant |B(r,s)|$ for all (r,s), then $|A(r,s)| = |B(r,s)|$ for all (r,s).*

The next example provides some further intuition for the above results.

[7] To be clear, we do not know whether such an allocation rule actually exists.

[8] Indeed, an AAR and a TU-maximizer yield the same number of value-creating matches when the change in marginal cost at the total utility maximizing output level is sufficiently big, *e.g.*, when the producer operates close to or at capacity.

Example 3.4.7 Consider a market with two consumers and one producer. Recall that the reported valuations and marginal costs are, respectively, given by r_1, r_2, s_1 and s_2, with $s_1 \leqslant s_2$. Let σ be a permutation such that $r_{\sigma(1)} \geqslant r_{\sigma(2)}$. Figure 3.2 depicts the situation under a TU-maximizer, whereas the situation under an AAR is illustrated in Figure 3.3. Note that the area in which two items are consumed is smaller in case of total utility maximization (comparing Figure 3.2(a) and Figure 3.3(a)). Similarly, the area where the supplier produces two items is also smaller when maximizing total utility (comparing Figure 3.2(b) and Figure 3.3(b)).

Next, consider Figure 3.2(b), which depicts the situation for the producer under a TU-maximizer. Note that value-creation requires that the vertical line marking the boundary between the area with no item and one item remains unchanged. Moreover, monotonicity requires the boundary between the area where the producer produces one and the area where it produces two items to be a straight horizontal line. Hence, when an implementable, value-creating allocation rule yields weakly more items than a TU-maximizer, the horizontal line at $r_{\sigma(2)}$ shifts upward. In order to maintain value-creation, it can maximally shift to $\frac{r_{\sigma(1)}+r_{\sigma(2)}}{2}$, which is precisely what an Average Allocation Rule does. This expansion can also be seen directly in Figure 3.3, where the gray region captures the additional set of valuations, relative to a TU-maximizer, for which the allocation shifts from one to two units. Consider point A in Figure 3.3(a). At this valuation profile, the AAR allocates two units, whereas a TU-maximizer allocates only one unit to consumer 1. Similarly, point B in Figure 3.3(b) illustrates a producer cost profile for which the AAR yields the production of two units, while a TU-maximizer yields the production of only one unit.

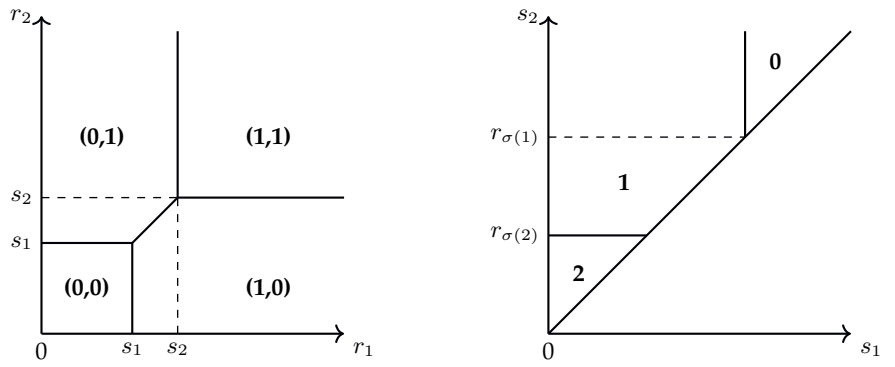

(a) Consumers (b) Producer

Figure 3.2: An illustration of a TU-maximizer in Example 3.4.7.

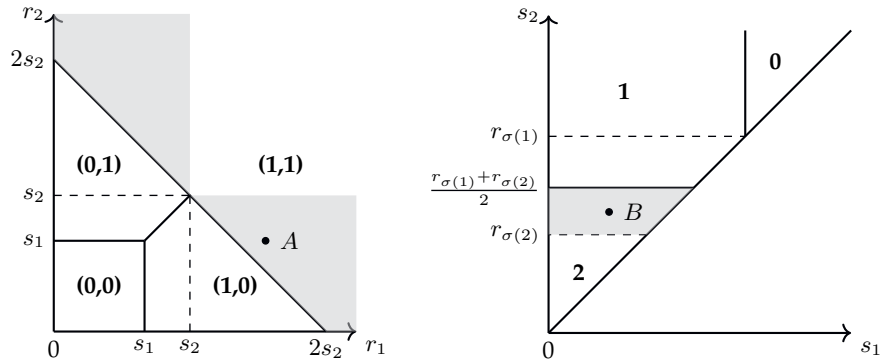

(a) Consumers (b) Producer

Figure 3.3: An illustration of an AAR in Example 3.4.7.

3.5 Bilateral Value-Creation

In the previous section, we presented a rule (an AAR) that dominates the utilitarian rule in terms of the number of transactions when marginal costs are non-decreasing. Note, however, that such a rule is value-creating (Definition 3.2.1), but not always bilaterally value-creating (Definition 3.2.3). In practice, it is not always possible to cross-subsidize transactions and therefore it may be necessary for matchings to be bilaterally value-creating (Definition 3.2.2). In this case, it appears that the utilitarian rule is again undominated.

We maintain the assumption that marginal costs are non-decreasing. The next theorem shows that any monotonic and value-creating allocation rule that yields at least the number of transactions under TU maximization also yields at most the number of transactions under TU maximization.

Theorem 3.5.1 *Let A be any monotonic and bilaterally value-creating allocation rule and suppose that A^{TU} is a TU-maximizer with maximal cardinality. If $|A(r,s)| \geqslant |A^{TU}(r,s)|$ for all (r,s), then $|A(r,s)| = |A^{TU}(r,s)|$ for all (r,s).*

The proof proceeds by constructing alternative cost profiles that make the additional matches chosen by A unattractive, while preserving the number of matches under A^{TU}. If A were to maintain a weakly larger number of matches, this would either violate bilateral value-creation (by assigning inefficient pairs) or monotonicity (by reducing the number of matches when costs decrease). By carefully distinguishing cases depending on the marginal costs, the argument ensures that any deviation from A^{TU} leads to such a contradiction. Hence, A cannot outperform A^{TU} in cardinality.

3.6 Multi-Unit Demand

The preceding results were derived under the assumption of unit demand. In this section, we take a step toward exploring the possibility of multi-unit demand. Specifically, we consider a situation in which there is one consumer and one producer.

Suppose that the consumer can demand $m \geqslant 1$ items.[9] The reported marginal valuation for the kth item is given by $r_k \geqslant 0$. It is assumed that $r_1 \geqslant r_2 \geqslant \ldots \geqslant r_m$ and the set of all such reports is indicated by R. In a similar vein, let $s_k \geqslant 0$ be the reported marginal cost of the kth item. We assume that $s_1 \leqslant s_2 \leqslant \ldots \leqslant s_m$ and denote the set of all such reports by S. An allocation rule assigns a number $q(r,s) \in \mathbb{N}$ for all $(r,s) \in R \times S$.

To study this setting, we restate Definition 3.2.1 and Definition 3.2.6 as follows.

Definition 3.6.1 *An allocation rule q is **value-creating** if, for all $r \in R$ and $s \in S$, it holds that:*

$$\sum_{i=1}^{k} r_i \geqslant \sum_{i=1}^{k} s_k \quad \text{where } k = q(r,s).$$

Definition 3.6.2 *An allocation rule q is **monotonic** if, for all $r, r' \in R$ and $s, s' \in S$, it holds that*

$$\sum_{i=l+1}^{k} r_i \geqslant \sum_{i=l+1}^{k} r'_i \quad \text{and} \quad \sum_{i=j+1}^{k} s_i \leqslant \sum_{i=j+1}^{k} s'_i$$

where $k = q(r,s)$, $l = q(r',s)$, and $j = q(r,s')$ with $l < k$ and $j < k$.

The next theorem establishes that in this case maximizing total utility is equivalent to maximizing the number of transactions.

Theorem 3.6.3 *Let q be any monotonic and value-creating allocation rule and suppose that q^{TU} is a TU-maximizer with maximal cardinality. If $q(r,s) \geqslant q^{TU}(r,s)$ for all $r \in R$ and $s \in S$, then $q(r,s) \leqslant q^{TU}(r,s)$ for all $r \in R$ and $s \in S$.*

The key idea is to construct alternative valuation and cost profiles that make the additional matches in q unattractive while keeping the TU-maximizing allocation unchanged. Under these profiles, the extra matches violate value-creation and cause q to reduce the number of matches even though the relevant costs decrease, contradicting monotonicity. By carefully adjusting the marginal valuations and costs in several cases, the argument ensures that any allocation exceeding q^{TU} leads to such a contradiction, showing that q must coincide with q^{TU}.

[9] Myerson and Satterthwaite (1983) provide a detailed study of a situation in which one buyer attempts to purchase a single item from one seller.

3.7 Concluding Remarks

In this chapter, we have examined allocation rules that maximize trade volume in monopoly. We introduced the concepts of value-creation and bilateral value-creation to formalize feasibility constraints. Our results show that, without any particular assumption on production costs, no implementable and value-creating allocation rule can consistently facilitate more transactions than total utility maximization. In case marginal costs are non-decreasing, however, there is an allocation rule (which we coined Average Allocation Rule) that yields weakly more transactions while preserving implementability and value-creation. Yet, this rule does not satisfy bilateral value-creation. We prove that there is no implementable, bilaterally value-creating allocation rule that can outperform the utilitarian rule in terms of the number of transactions.

There are two natural avenues for further investigation. One is to explore whether our bilateral monopoly result is robust to allowing for increasing marginal valuations. The other is to study trade maximizing allocation rules in a non-monopolistic setting.

Appendix: Proofs

Proof of Theorem 3.3.3

Take any monotonic and value-creating allocation rule A with $|A(r,s)| \geqslant |A^{TU}(r,s)|$ for all r, s. We prove that $|A^{TU}(r,s)| \geqslant |A(r,s)|$ for all r, s.

By contradiction, assume that there is a profile (r,s) with $|A^{TU}(r,s)| = n < k = |A(r,s)|$. Since we keep r fixed throughout the proof, we can assume without loss of generality that $r_1 \geqslant \cdots \geqslant r_m$. We construct a cost profile s' with $|A(r,s')| = k' < k$ and $\sum_{i=k'+1}^{k} s'_i < \sum_{i=k'+1}^{k} s_i$. By the remark following Definition 3.2.6 this contradicts the monotonicity of A.

Since $A^{TU}(r,s) = n < k$, and A^{TU} is a maximum cardinality TU maximizer, we know that

$$TU(\{1,\ldots,n\})(r,s) = TU(A(r,s))(r,s) > TU(\{1,\ldots,k\})(r,s).$$

This is equivalent to $\sum_{i=n+1}^{k} s_i > \sum_{i=n+1}^{k} r_i$. Write $\varepsilon = \sum_{i=n+1}^{k} s_i - \sum_{i=n+1}^{k} r_i$. Then $\varepsilon > 0$. Write $\mu = \sum_{i=1}^{n} r_i - \sum_{i=1}^{n} s_i$. Define the costs profile $\mathbf{s'} = (s'_1, \ldots, s'_m)$ by

$$s'_i = \begin{cases} s_1 + \mu & \text{if } i = 1 \\ s_i & \text{if } 2 \leqslant i \leqslant k-1 \\ s_k - \frac{\varepsilon}{2} & \text{if } i = k \\ m \cdot K & \text{if } k+1 \leqslant i \leqslant m, \end{cases} \qquad \text{where } K = \max\{r_i \mid i \in M\} + 1.$$

A. We show that $|A(r,s')| < k$. Recall that $\varepsilon > 0$. So,

$$\begin{aligned} \sum_{i=1}^{k} s'_i = \sum_{i=1}^{k} s_i + \mu - \frac{\varepsilon}{2} &= \sum_{i=1}^{k} s_i + \sum_{i=1}^{n} r_i - \sum_{i=1}^{n} s_i - \frac{\varepsilon}{2} \\ &= \sum_{i=n+1}^{k} s_i + \sum_{i=1}^{k} r_i - \sum_{i=n+1}^{k} r_i - \frac{\varepsilon}{2} \\ &= \sum_{i=1}^{k} r_i + \frac{\varepsilon}{2} > \sum_{i=1}^{k} r_i. \end{aligned}$$

So, if $|A(r,s')| = k$, then

$$\sum_{i=1}^{k} s'_i > \sum_{i=1}^{k} r_i \geqslant \sum_{i \in A(r,s')} r_i.$$

This violates the assumption that A is value-creating. Further, if $|A(r,s')| > k$, then, since $s'_{k+1} = m \cdot K$, it holds that

$$\sum_{i=1}^{k'} s'_i \geqslant m \cdot K > \sum_{i=1}^{m} r_i \geqslant \sum_{i \in A(r,s')} r_i.$$

Again, this violates the assumption that A is value-creating. Hence, $|A(r, s')| < k$.

B. We show that

$$\sum_{i=k'+1}^{k} s'_i < \sum_{i=k'+1}^{k} s_i.$$

We consider two cases.

B1. If $n = 0$. Then, since A^{TU} is a maximum cardinality total utility maximizer,

$$\sum_{i=1}^{\ell} s_i > \sum_{i=1}^{\ell} r_i$$

for all $\ell = 1, \ldots, m$. Then, since A is value-creating, also $k = 0$. So, this is excluded by our assumption that $n < k$.

B2. If $n \geqslant 1$. Note that

$$\sum_{i=1}^{n} s'_i = \sum_{i=1}^{n} s_i + \mu = \sum_{i=1}^{n} r_i.$$

So, allocating n items has the same total utility as allocating zero items. Then, since A^{TU} is a maximum cardinality total utility maximizer, $|A^{TU}(r, s')| \geqslant 1$. Then also $k' = |A(r, s')| \geqslant 1$. So, since $s'_k = s_k - \frac{\varepsilon}{2}$ and $s'_i = s_i$ for all $2 \leqslant i < k$, it holds that

$$\sum_{i=k'+1}^{k} s'_i = \sum_{i=k'+1}^{n} s_i - \frac{\varepsilon}{2} < \sum_{i=k'+1}^{k} s_i.$$

Thus, according to Definition 3.2.6, A is not monotonic. ■

The next result is used in the proof of Proposition 3.4.3.

Lemma A.1 *Let σ be a permutation such that $r_{\sigma(1)} \geqslant \cdots \geqslant r_{\sigma(m)}$. Then,*

$$\frac{1}{k} \cdot \sum_{i=1}^{k} r_{\sigma(i)} \geqslant \frac{1}{k+1} \cdot \sum_{i=1}^{k+1} r_{\sigma(i)} \quad \textit{for all } k = 1, \ldots, m-1.$$

Proof of Lemma A.1

Suppose, without loss of generality, that σ is the identity permutation. So, we know that $r_1 \geqslant \ldots \geqslant r_n$ and, therefore, that $\frac{r_1}{1} = \frac{r_1+r_1}{2} \geqslant \frac{r_1+r_2}{2} \geqslant \frac{r_1+r_2+r_3}{3} \geqslant \ldots$. ■

Proof of Proposition 3.4.3

Let A be an AAR. We first prove that A is value-creating. Suppose that $|A(r, s)| = k$ and let σ be a permutation with $r_{\sigma(1)} \geqslant \cdots \geqslant r_{\sigma(m)}$. Then, $\frac{1}{k} \cdot \sum_{i=1}^{k} r_{\sigma(i)} \geqslant s_k$. Hence,

since $s_k \geqslant \cdots \geqslant s_1$, it follows that:

$$\sum_{i \in A(r,s)} r_i = \sum_{i=1}^{k} r_{\sigma(i)} \geqslant k \cdot s_k \geqslant \sum_{i=1}^{k} s_i = s(k).$$

Next we prove that A is monotonic. Consider the jth consumer and let (r, s) be a report. Suppose that $j \in A(r, s)$. Take any $r'_j \geqslant r_j$. For $r' = (r'_j, r_{-j})$, we have to show that $j \in A(r', s)$.

Write $k = |A(r, s)|$ and $k' = |A(r', s)|$. Suppose, without loss of generality, that $r_1 \geqslant \cdots \geqslant r_m$ and that τ is such that $r'_{\tau(1)} \geqslant \cdots \geqslant r'_{\tau(m)}$.

Since $r'_j \geqslant r_j$, we know that $\ell \leqslant j$ for ℓ with $\tau(\ell) = j$. Since with $j \in A(r, s)$, we obtain

$$s_k \leqslant \frac{1}{k} \cdot \sum_{i \in A(r,s)} r_i \leqslant \frac{1}{k} \cdot \left(\sum_{\substack{i \in A(r,s) \\ i \neq j}} r_i + r'_j \right).$$

It follows that $k' \geqslant k$. Hence, since $\ell \leqslant j \leqslant k \leqslant k'$, it follows that $i \in A(r', s)$.

Consider the producer. Let r be any report by the consumers and let s and s' be any two reports by the producer. Write $k = |A(r, s)|$ and $k' = |A(r, s')|$. Without loss of generality, assume that $k' > k$.

We have to prove that:

$$\sum_{i=k+1}^{k'} s'_i \leqslant \sum_{i=k+1}^{k'} s_i.$$

We show the stronger statement that $s'_\ell < s_\ell$, for all $\ell = k + 1, \ldots, k'$. Take a fixed $\ell = k + 1, \ldots, k'$. Since $k' = |A(r, s')|$ and $\ell \leqslant k'$, we know that

$$s'_\ell \leqslant s'_{k'} \leqslant \frac{1}{k'} \cdot \sum_{i=1}^{k'} r_i.$$

Since $k = |A(r, s')|$, and $\ell > k$, we know that

$$\frac{1}{k+1} \cdot \sum_{i=1}^{k+1} r_i < s_{k+1} \leqslant s_\ell.$$

Hence, since $k + 1 \leqslant k'$, by Lemma A.1, we have that

$$s'_\ell \leqslant \frac{1}{k'} \cdot \sum_{i=1}^{k'} r_i \leqslant \frac{1}{k+1} \cdot \sum_{i=1}^{k+1} r_i < s_\ell.$$

This completes the proof. ■

The next Lemma is used in the proof of Theorem 3.4.4.

Lemma A.2 *Without loss of generality, assume that consumers are indexed such that* $r_1 \geqslant \cdots \geqslant r_m$. *Under any TU-maximizer,* $r_i \geqslant s_i$ *for all* $i \leqslant |A^{TU}(r,s)|$.

Proof of Lemma A.2

Since a TU-maximizer allocates $|A^{TU}(r,s)| = k$ items in total, we know that $\sum_{i=1}^{k} r_i - \sum_{i=1}^{k} s_i \geqslant \sum_{i=1}^{k-1} r_i - \sum_{i=1}^{k-1} s_i$. Hence, $r_k \geqslant s_k$. Since consumers are indexed such that $r_1 \geqslant \ldots \geqslant r_k$ and marginal costs are non-decreasing, it follows that $r_1 \geqslant \ldots \geqslant r_k \geqslant s_k \geqslant \ldots \geqslant s_1$. ■

Proof of Theorem 3.4.4

Let (r,s) be any report. Write $k = |A^{TU}(r,s)|$, and without loss of generality, assume that $r_1 \geqslant \cdots \geqslant r_m$. Since a TU-maximizer allocates k items for the producer, we know that $r_k \geqslant s_k$ by Lemma A.2. Hence,

$$\frac{1}{k} \cdot \sum_{i=1}^{k} r_i \geqslant \frac{k.r_k}{k} = r_k \geqslant s_k.$$

Therefore, $|A(r,s)| \geqslant k$. ■

Proof of Theorem 3.4.6

Let A be any AAR rule, and let B be any monotonic and value-creating allocation rule. Suppose that $|A(r,s)| \leqslant |B(r,s)|$ for every report (r,s). Take any report (r,s). We prove that $|B(r,s)| = |A(r,s)|$.

Assume that $|A(r,s)| < n = |B(r,s)|$. We derive a contradiction with monotonicity of B as in Definition 3.2.6. We construct a cost report s' with $|B(r,s')| = n-1$ but $s'_n < s_n$.

Since we only change costs, we may assume without loss of generality that $r_1 \geqslant \cdots \geqslant r_m$. Then, since $|A(r,s)| < n$, we know that $\frac{1}{n} \cdot \sum_{i=1}^{n} r_i < s_n$.

A. If $n = 1$. Then $|A(r,s)| = 0$. So, $s_j > r_i$ for any $i, j \in M$. Then also $|B(r,s)| = 0$ due to value-creation. This contradicts the assumption that $|A(r,s)| < |B(r,s)|$.

B. So, $n \geqslant 2$. Write $\varepsilon = s_n - \frac{1}{n} \cdot \sum_{i=1}^{n} r_i$. Clearly $\varepsilon > 0$. Define $s' = (s'_1, s'_2, \ldots)$ by

$$s'_i = \begin{cases} \frac{1}{n} \cdot \sum_{i=1}^{n} r_i & \text{if } 1 \leqslant i \leqslant n-1 \\ \frac{1}{n} \cdot \sum_{i=1}^{n} r_i + \frac{\varepsilon}{2} & \text{otherwise.} \end{cases}$$

Notice that indeed s' has non-decreasing marginal cost. Since $r_1 \geqslant \cdots \geqslant r_m$, we obtain

$$\frac{\sum_{i=1}^{n-1} r_i}{n-1} \geqslant \frac{\sum_{i=1}^{n} r_i}{n} = s'_{n-1}, \text{ and}$$
$$\frac{\sum_{i=1}^{n} r_i}{n} < \frac{\sum_{i=1}^{n} r_i}{n} + \frac{\varepsilon}{2} = s'_n.$$

Hence, $|A(r,s')| = n-1$. So, since $|B(r,s')| \geqslant |A(r,s')|$, we have that $|B(r,s')| \geqslant n-1$. Then, since

$$\sum_{i \in B(r,s')} r_i \leqslant \sum_{i=1}^{n} r_i < \sum_{i=1}^{n} r_i + \frac{\varepsilon}{2} = \sum_{i=1}^{n} s'_i$$

and B is value-creating, it follows that $|B(r,s')| = n-1$. Further,

$$s'_n = \frac{\sum_{i=1}^{n} r_i}{n} + \frac{\varepsilon}{2} < \frac{\sum_{i=1}^{n} r_i}{n} + \varepsilon = s_n.$$

So, according to Definition 3.2.6, B is not monotonic. This contradicts the assumption that B is monotonic. ■

The next Lemma is used in the proof of Theorem 3.5.1.

Lemma A.3 *Consider a bilaterally value-creating matching T that yields k matches. Let K denote the set of consumers matched under T. Write $K = \{1,\ldots,k\}$ such that $r_i \geqslant r_{i+1}$. Consider the matching $T' = \{(i, k+1-i) \mid 1 \leqslant i \leqslant k\}$. Then T' is bilaterally value-creating.*

Proof of Lemma A.3

Consider consumer 1 and item k. Let j' denote the item consumer 1 gets under T, and let i' denote the consumer who gets item k under T.

If $i' = 1$ (and hence $k = j'$), the claim follows immediately from the bilateral value-creating of T, which implies $r_1 \geqslant s_k$. If $i' \neq 1$, since consumer 1 has the highest valuation among all consumers in K, we have $r_1 \geqslant r_{i'}$. Because (i', k) is a match in T, bilateral value-creating implies $r_{i'} \geqslant s_k$. Hence, $r_1 \geqslant s_k$. Moreover, since k is the highest marginal-cost item that is produced, we have $s_k \geqslant s_{j'}$. Combining this with $r_{i'} \geqslant s_k$ yields $r_{i'} \geqslant s_{j'}$.

We can therefore replace the pairs $(1, j')$ and (i', k) with $(1, k)$ and (i', j') without violating bilateral value-creating. Remove consumer 1 from K, update the matching accordingly, and repeat this argument for the remaining agents. In the end, we obtain a bilaterally value-creating matching $T' = \{(i, k+1-i) \mid 1 \leqslant i \leqslant k\}$. ■

Proof of Theorem 3.5.1

Take any monotonic and bilaterally value-creating allocation rule A with $|A(r,s)| \geqslant |A^{TU}(r,s)|$ for all reports (r,s). We show that $|A(r,s)| \leqslant |A^{TU}(r,s)|$ for all reports (r,s).

By contradiction, assume that there is a report (r,s) with $|A^{TU}(r,s)| = k < n = |A(r,s)|$. Since we keep r fixed throughout the proof, we can assume without loss of generality that $r_1 \geqslant \cdots \geqslant r_m$. We construct a costs profile s' with $\ell = |A(r,s')| < n$ but $\sum_{i=\ell+1}^{n} s'_i < \sum_{i=\ell+1}^{n} s_i$. This contradicts the monotonicity of A.

A. Assume that $r_k > r_{k+1}$. Since $|A^{TU}(r,s)| = k$, we know that $r_{k+1} < s_{k+1}$. Choose $\varepsilon > 0$ with

$$r_k > s_{k+1} - \varepsilon > r_{k+1}.$$

Define reported costs profile $s' = (s'_1, \ldots, s'_m)$ by

$$s'_i = \begin{cases} s_{k+1} - \varepsilon & \text{if} \quad 1 \leqslant i \leqslant n \\ r_1 + 1 & \text{if } n+1 \leqslant i \leqslant m. \end{cases}$$

Note that s' has non-decreasing marginal cost. We show that $|A^{TU}(\mathbf{r}, \mathbf{s}')| = k$. By choice of ε, we have $r_k > s_{k+1} - \varepsilon = s'_k$. So, $|A^{TU}(r,s')| \geqslant k$. Further, $r_{k+1} < s_{k+1} - \varepsilon = s'_{k+1}$ by choice of ε. So, $|A^{TU}(r,s')| < k+1$. It follows that $|A^{TU}(r,s')| = k$.

Write $|A(\mathbf{r}, \mathbf{s}')| = \ell$. Then $\ell \geqslant k$ by assumption. Further, $\ell \leqslant n$ by bilateral rationality of A. We first argue that $\ell = n$. Suppose that $\ell < n$. Since $\ell \geqslant k$, it holds that

$$\sum_{i=\ell+1}^{n} s'_i = \sum_{i=\ell+1}^{n} s_{k+1} - (n - \ell - 1) \cdot \varepsilon < \sum_{i=\ell+1}^{n} s_{k+1} \leqslant \sum_{i=\ell+1}^{n} s_i.$$

So, according to Definition 3.2.6, A is not monotonic. So, $\ell \geqslant n$. However, since $s'_i = r_1 + 1 > r_1 \geqslant \ldots \geqslant r_m$ for all $n+1 \leqslant i \leqslant m$, and A is bilaterally value-creating, $\ell < n+1$. Hence, $\ell = n$.

We derive a contradiction. Since $\ell = n$, we know by Lemma A.3 that for all $\ell \geqslant n$, $r_n \geqslant r_\ell \geqslant s'_1 = s_{k+1} - \varepsilon > r_{k+1}$. Since $k < n$, this is a contradiction.

B1. Assume that $r_k = r_{k+1}$. We construct a cost profile s' with $|A^{TU}(r,s')| = k$ and $|A(r,s')| = n$. We know that $s_{k+1} > r_{k+1} = r_k$. Take $\varepsilon > 0$ such that $s_{k+1} - \varepsilon \geqslant r_k$. Define $s' = (s'_1, \ldots, s'_m)$ by

$$s'_i = \begin{cases} r_k & \text{for all } 1 \leqslant i \leqslant k, \\ r_k + \frac{\varepsilon}{2} & \text{for all } k+1 \leqslant i \leqslant n \\ r_1 + 1 & \text{for all } n+1 \leqslant i \leqslant m. \end{cases}$$

We show that $|A^{TU}(r,s')| = k$. Note that marginal costs are non-decreasing. Because $s'_k = r_k$ and $r_{k+1} = r_k < r_k + \frac{\varepsilon}{2} = s'_{k+1}$, we know that $|A^{TU}(r,s')| = k$.

Write $|A(r,s')| = \ell$. We show that $\ell = n$. If $\ell < n$. Note that $\ell \geqslant k$ by assumption. So, by choice of ε,

$$\sum_{i=\ell+1}^{n} s'_i = \sum_{i=\ell+1}^{n} (r_k + \frac{\varepsilon}{2}) < \sum_{i=\ell+1}^{n} s_{k+1} \leqslant \sum_{i=\ell+1}^{n} s_i.$$

So, according to Definition 3.2.6, A is not monotonic. This is a contradiction. So, $\ell \geqslant n$. Then, since $r_1 < r_1 + 1 = s_{n+1}$, we have $\ell \leqslant n$ by bilateral rationality.

B2. We derive a contradiction by choosing a costs profile s'' such that for $\ell' = |A(r,s'')|$ we have $\ell' < n$ and $\sum_{i=\ell'+1}^{n} s''_i < \sum_{i=\ell'+1}^{n} s'_i$. Define $s'' = (s''_1, \ldots, s''_m)$ by

$$s''_i = \begin{cases} r_k & \text{for all } 1 \leqslant i \leqslant n-k \\ r_k + \frac{\varepsilon}{4n} & \text{for all } n-k+1 \leqslant i \leqslant n \\ r_1 + 1 & \text{for all } n+1 \leqslant i \leqslant m. \end{cases}$$

Note that s'' is non-decreasing. Write $|A(r,s'')| = \ell'$. We show that $\ell' < n$. Since $s''_{n-k+1} = r_k + \frac{\epsilon}{4n} > r_k$, consumer k cannot get the $(n-k+1)^{st}$ item from the producer. Hence, according to Lemma A.3, we obtain that $\ell' < n$. Then we have

$$\sum_{i=\ell'+1}^{n} s''_i \leqslant (n-\ell') \cdot \left(r_k + \frac{\varepsilon}{4n}\right) \leqslant (n-\ell') \cdot r_k + \frac{\varepsilon}{4} < \sum_{i=\ell'+1}^{n} s'_i.$$

According to Definition 3.2.6, A is not monotonic. This is a contradiction. ■

Proof of Theorem 3.6.3

Consider a monotonic and value-creating allocation rule q with $q(r,s) \geqslant q^{TU}(r,s)$ for all $r \in R$ and $s \in S$. We prove that $q(r,s) \leqslant q^{TU}(r,s)$ for all $r \in R$ and $s \in S$. Assume that for some profile (r,s) we have that $q^{TU}(r,s) < q(r,s)$. We derive a contradiction.

Write $n = q(r,s)$. Since $q^{TU}(r,s) < n$, we know that $s_n > r_n$. Write $\varepsilon = \frac{s_n - r_n}{n+1}$. Then $\varepsilon > 0$. Define reported valuations profile $r' = (r'_1, \ldots, r'_m)$ by

$$r'_i = \begin{cases} s_n & \text{if } i = 1, \ldots, n-1 \\ s_n - \frac{\varepsilon}{2} & \text{if } i = n \\ 0 & \text{if } i = n+1, \ldots, m. \end{cases}$$

Clearly $r' \in R$.

Claim A. It holds that $q^{TU}(r',s) = n-1$ and $q(r',s) \geqslant n$.

Proof of Claim A. Note that

$$r'_{n-1} = s_n \geqslant s_{n-1} \quad \text{and} \quad r'_n = s_n - \frac{\varepsilon}{2} < s_n.$$

So, $q^{TU}(r', s) = n - 1$. It follows that $q(r', s) \geqslant n - 1$. Then, since

$$r'_n = s_n - \frac{\varepsilon}{2} > s_n - (n+1) \cdot \varepsilon = r_n,$$

it follows from Definition 3.6.2 that $q(r', s) \geqslant n$. End proof of Claim A.

We derive a contradiction in two steps, depending on whether $q(r', s) = n$ or $q(r', s) > n$.

If $q(r', s) = n$. Define reported costs profile $s' = (s'_1, \ldots, s'_m)$ by

$$s'_i = \begin{cases} s_n - \frac{\varepsilon}{3 \cdot n} & \text{if } i = 1, \ldots, n-1 \\ s_n - \frac{\varepsilon}{6 \cdot n} & \text{if } i = n \\ s_n + 1 & \text{if } i = n+1, \ldots, m. \end{cases}$$

Clearly $s' \in S$.

Claim B. It holds that $q^{TU}(r', s') = n - 1$ and $q(r', s') = n - 1$.

Proof of Claim B. Note that

$$s'_{n-1} = s_n - \frac{\varepsilon}{3 \cdot n} \leqslant s_n = r'_{n-1} \quad \text{and} \quad s'_n = s_n - \frac{\varepsilon}{6 \cdot n} > s_n - \frac{\varepsilon}{2} = r'_n.$$

Then, $q^{TU}(r', s') = n - 1$. Then also $q(r', s') \geqslant n - 1$. So, since

$$\sum_{i=1}^{n} s'_i = n \cdot s_n - \frac{\varepsilon}{3} + \frac{\varepsilon}{6 \cdot n} > \sum_{i=1}^{n} r'_i$$

we have that $q(r', s') = n - 1$ by value-creation of q. End proof of Claim B.

Now note that

$$s'_n = s_n - \frac{\varepsilon}{6 \cdot n} < s_n.$$

So, q is not monotonic by Definition 3.6.2 and Claims A and B. This concludes the case where $q(r', s) = n$.

If $q(r', s) = n + k$, with $k \geqslant 1$. Define $s'' = (s''_1, \ldots, s''_m)$ by

$$s''_i = \begin{cases} s_n - \varepsilon & \text{if } i = 1, \ldots, n+1 \\ s_i & \text{if } i = n+2, \ldots, m. \end{cases}$$

Clearly $s'' \in S$.

Claim C. It holds that $q^{TU}(r', s'') = n - 1$ and $q(r', s'') \in \{n - 1, n\}$.

Proof of Claim C. Further,

$$s''_n = s_n - \varepsilon < s_n - \frac{\varepsilon}{2} = r'_n \quad \text{and} \quad s''_{n+1} = s_n - \varepsilon = \frac{n \cdot s_n + r_n}{n+1} > 0 = r'_{n+1}.$$

So, $q^{TU}(r', s'') = n - 1$. Then also $q(r', s'') \geqslant n - 1$. So, since

$$\sum_{i=1}^{n+1} s''_i = (n+1) \cdot (s_n - \varepsilon) = n \cdot s_n + r_n > n \cdot s_n - \frac{\varepsilon}{2} = \sum_{i=1}^{n+1} r'_i,$$

it follows from value-creation of q that $n - 1 \leqslant q(r', s'') \leqslant n$. However, we have $s''_n = s''_{n+1} = s_n - \varepsilon < s_n \leqslant s_{n+1}$ and $s'_i = s_i$ for all $i = n+2, \ldots, m$. It follows that

$$\sum_{i=n+j}^{n+k} s''_i < \sum_{i=n+j}^{n+k} s_i, \quad \text{for all } j \in \{0, 1\}.$$

Hence, by Definition 3.6.2, q is not monotonic. This is a contradiction. ■

4

Characterizing Total Utility Maximization in Matching Markets with Heterogeneous Valuations

This chapter is based on joint work with Iwan Bos, Marc Schröder and Dries Vermeulen (Maastricht University).

4.1 Introduction

Allocating indivisible objects among agents with heterogeneous valuations is a central problem in a wide range of economic settings, including procurement, frequency allocation, and school assignment. Each object can be assigned to at most one agent, and agents may value objects differently. A natural objective is to maximize total utility, selecting a matching that maximizes aggregate welfare. In the absence of transfers, maximizing total utility reduces to finding a maximum weight matching in a bipartite graph (see, *e.g.*, Kuhn (1955)). When transfers are allowed, it corresponds to the surplus-maximizing allocation in the assignment game and underlies competitive equilibrium outcomes (L. S. Shapley & Shubik, 1971).

Despite its appeal, total utility maximization is difficult to implement when agents' valuations are privately known. In single-dimensional environments, where each agent's type is described by a single parameter, mechanisms that induce truthful reports are characterized by the classical monotonicity condition on the allocation rule (Myerson, 1981). In multidimensional settings, however, as in the environments with heterogeneous valuations that we consider, each agent's type consists of a vector of valuations over multiple objects, and the design of incentive-compatible allocation rules becomes more challenging. Several results in the literature (see, *e.g.*, Bikhchandani et al. (2003), Lavi et al. (2003), and Rochet (1987)) establish that implementability in multidimensional domains is equivalent to monotonicity. While this equivalence provides a necessary and sufficient condition for implementability, it leaves an important open question: how restrictive are implementable allocation rules relative to the total utility benchmark? In particular, to what extent can an allocation rule satisfying monotonicity and other natural economic properties deviate from total utility maximization?

This chapter extends the analysis of Chapters 2 and 3, which study homogeneous-good markets with multiple unit-supply sellers and with a monopolist, respectively, to the case of heterogeneous products, where agents may value distinct objects differently. In this chapter, we consider allocation problems with heterogeneous cardinal valuations and at most one object per agent. Within this setting, we identify conditions under which an implementable rule cannot strictly dominate a maximal-cardinality total utility maximizer in terms of the number of matches. To this end, we introduce three natural economic properties: monotonicity, which ensures implementability in multidimensional settings; value-creation, which requires that objects are allocated only when they generate nonnegative utility; and irrelevance of unassigned objects, which ensures that lowering an agent's valuation for an object they do not receive does not affect the overall allocation. First, we show that any total utility-maximizing allocation rule satisfies all these properties. Then, we analyze allocation rules that satisfy monotonicity, value-creation, and irrelevance of unassigned objects. Our main result establishes a cardinality restriction: if an allocation rule satisfies these properties and weakly dominates a maximal-cardinality total utility maximizer in terms of matching size, then the two rules must coincide in cardinality for every valuation profile.

Our setting is related to several classical models of assignment and matching. When the number of agents equals the number of objects, it is closely related to the house allocation framework of (Hylland & Zeckhauser, 1979). Much of this literature studies efficiency and strategy-proofness in environments without transfers, often under ordinal preferences (see, *e.g.*, Abdulkadiroğlu and Sönmez (1998), Ehlers (2002), and Ehlers et al. (2002)). The introduction of cardinal utilities in such settings leads to well-known impossibility results: Zhou (1990) confirms a conjecture of Gale and Shapley (1962), showing that no mechanism can simultaneously satisfy strategy-proofness, ex-ante efficiency, and symmetry. More recent work addresses this tension through randomization. For example, Abebe et al. (2020) propose mechanisms that elicit truthful cardinal reports and achieve approximate efficiency.

In contrast, we consider cardinal environments with transfers, corresponding to the model of Demange et al. (1986) and closely related to the assignment game of L. S. Shapley and Shubik (1971). While this literature focuses on auction dynamics and equilibrium price adjustments, we instead analyze how incentive compatibility, captured by monotonicity, together with other natural economic properties, constrains the size of feasible matchings.

The remainder of the chapter is organized as follows. Section 4.2 introduces the model, notation, and formal definitions of the key properties. Section 4.3 presents the main result. Section 4.4 concludes. The proofs of all auxiliary lemmas are provided in the Appendix.

4.2 Preliminaries

4.2.1 Model

Consider an allocation problem with a set of agents $M = \{1, \dots, m\}$ and a set of n indivisible objects to be allocated $N = \{o_1, \dots, o_n\}$. Each agent $i \in M$ demands at most one object and has heterogeneous valuations across objects. The valuation of agent i for object o is denoted by $v_{i,o} \in \mathbb{R}$, representing the utility that agent i derives from receiving o.

$T \subseteq M \times N$ is a matching if for every $i \in M$, $|\{o \mid (i,o) \in T\}| \leqslant 1$, and for every $o \in N$, $|\{i \mid (i,o) \in T\}| \leqslant 1$. Denote the set of all feasible matchings as $\mathcal{T}$.

An allocation rule is a function

$$A \colon \mathbb{R}^{m \times n} \longrightarrow \mathcal{T}.$$

For a valuation profile $v = (v_{i,o})_{i \in M,\, o \in N}$, the matching selected by the allocation rule is denoted by $A(v) \in \mathcal{T}$.

We next formalize the notion of total utility maximization. Since multiple matchings may maximize total utility, we refine the selection by focusing on those with maximal cardinality.

For a valuation profile v, define

$$U^{TU}(v) = \max_{T \in \mathcal{T}} U(T, v), \text{ where } U(T, v) = \sum_{(i,o) \in T} v_{i,o},$$

and

$$k^{TU}(v) = \max\{|T| \mid T \in \mathcal{T} \text{ and } U(T, v) = U^{TU}(v)\}.$$

We denote by $\mathcal{T}^{TU}(v)$ the set of all matchings $T \in \mathcal{T}$ such that

$$U(T, v) = U^{TU}(v) \quad \text{and} \quad |T| = k^{TU}(v).$$

Definition 4.1 *An allocation rule A is a **maximal-cardinality TU-maximizer** if for every valuation profile v, $A(v) \in \mathcal{T}^{TU}(v)$.*

We now introduce two properties of allocation rules that will play an important role in the remainder of the analysis.

Definition 4.2 *An allocation rule A is **value-creating** if for every profile $(v_{i,o})_{i \in M,\, o \in N}$ and every $(i, o) \in A(v)$, it holds that $v_{i,o} \geqslant 0$.*

That is, an object is allocated to an agent only when the agent's valuation for that object is nonnegative.

Let $v_{-i} \in \mathbb{R}^{(m-1) \times n}$ denote the valuation profile of all agents other than agent i. Furthermore, let $v_{-(i,o)} \in \mathbb{R}^{mn-1}$ denote the collection of all valuations except agent i's valuation for object o.

Definition 4.3 *An allocation rule A satisfies **irrelevance of unassigned objects (IUO)** if for all $i \in M$, for all $v_i \in \mathbb{R}^m$, for all o such that $(i, o) \notin A(v)$, and for all $v'_{i,o} < v_{i,o}$,*

$$A(v'_i, v_{-i}) = A(v_i, v_{-i}).$$

That is, if agent i is not assigned object o, lowering agent i's reported valuation for o should not change the allocation.

Remark 4.4 *Consider a value-creating allocation rule A that satisfies IUO. Let v be a preference profile such that $v_{i,o} < 0$. Then for every $v'_{i,o} < 0$,*

$$A(v'_{i,o}, v_{-(i,o)}) = A(v).$$

Intuitively, since A is value-creating, object o is never assigned to agent i whenever $v_{i,o} < 0$. Hence, by IUO, changing the value of $v_{i,o}$ within the negative region does not affect the allocation.

4.2.2 Graph-Theoretic Concepts

We introduce several graph-theoretic concepts that will be used in the proof of our main result (Theorem 4.2).

Definition 4.5 *A **directed bipartite graph** is a graph $\mathcal{G} = (V, E)$, where $V = M \cup N$, and*

$$E \subseteq (M \times N) \cup (N \times M).$$

That is, every directed edge connects a vertex in M to a vertex in N, or vice versa.

For a directed graph $\mathcal{G} = (V, E)$ and a vertex $u \in V$, the **in-degree** of u is the number of edges pointing to u, and the **out-degree** of u is the number of edges leaving u.

Definition 4.6 *A path $\mathcal{P}$ in the directed graph $\mathcal{G}$ is a finite ordered list of edges*

$$\mathcal{P} = ((u_1, u_2), (u_2, u_3), \ldots, (u_{L-1}, u_L)),$$

such that

1. *$u_i \in V$ for all $i = 1, \ldots, L$,*
2. *all vertices $u_1, \ldots, u_L$ are distinct,*
3. *$(u_i, u_{i+1}) \in E$ for all $i = 1, \ldots, L-1$.*

Definition 4.7 *A path $\mathcal{P}$ is **maximal** if there exists no vertex $u \in V \backslash \{u_1, \ldots, u_L\}$ such that either $(u, u_1) \in E$ or $(u_L, u) \in E$.*

After introducing these concepts, we state the following lemma, which will be used in the proof of Theorem 4.2.

Lemma 4.8 *Consider a finite directed bipartite graph $\mathcal{G} = (V, E)$ with bipartition $V = M \cup N$. Suppose that every vertex has in-degree and out-degree at most one, and that the number of edges from M to N is strictly larger than the number of edges from N to M. Then there exists at least one maximal path that starts at a node $i \in M$ and ends at a node $o \in N$.*

4.2.3 Implementation

From the perspective of the social planner, given the assumption of incomplete information, the planner may wish to employ a **direct mechanism** to achieve the desired allocation objective.[1]

A direct mechanism is a pair (A, p), where

$$A\colon \mathbb{R}^{m\times n} \longrightarrow \mathcal{T}$$

is the allocation rule, and

$$p\colon \mathbb{R}^{m\times n} \longrightarrow \mathbb{R}^{m}$$

is the payment rule specifying the transfer made by each agent.

The social planner first asks agents to report their valuations. Let

$$r = (r_{i,o})_{i\in M, o\in N} \in \mathbb{R}^{m\times n}$$

denote the matrix of reported valuations from all agents, where $r_{i,o}$ is the reported value of agent i for object o. Denote the vector of reported valuations of agent i for all items as $r_i \in \mathbb{R}^n$. Denote the vector of reported valuations for all agents except for agent i as $r_{-i} \in \mathbb{R}^{(m-1)\times n}$.

For each agent i, define the assignment vector $a_i(r) = (a_{i,o}(r))_{o\in N}$, where

$$a_{i,o}(r) = \begin{cases} 1 & \text{if } (i,o) \in A(r), \\ 0 & \text{otherwise.} \end{cases}$$

We proceed by introducing the key concepts to guarantee truthful behavior.

Definition 4.9 *A mechanism (A, p) is **dominant strategy incentive compatible (DSIC)** if, for every agent $i \in M$, for every pair of reported valuation vectors $r_i, r_i' \in \mathbb{R}^n$, and for every report profile of the other agents $r_{-i} \in \mathbb{R}^{(m-1)\times n}$, it holds that*

$$r_i \cdot a_i(r_i, r_{-i}) - p_i(r_i, r_{-i}) \;\geqslant\; r_i \cdot a_i(r_i', r_{-i}) - p_i(r_i', r_{-i}).$$

This indicates that the net utility of reporting the true valuations is higher than or equal to the net utility of reporting anything else. Therefore, in a DSIC mechanism, it is the dominant strategy for agents to report truthfully.

Definition 4.10 *An allocation rule A is **implementable** if there is a payment rule p such that (A, p) is DSIC.*

[1] We invoke the revelation principle by Gibbard (1973), which states that for any equilibrium outcome of a general mechanism, there exists an equivalent incentive-compatible direct revelation mechanism.

Bikhchandani et al. (2003) showed that, in a multi-dimensional setting, an allocation rule is implementable precisely when it is monotonic. A similar result regarding the equivalence of implementability and monotonicity was obtained by Lavi et al. (2003), though their proof follows a different approach.

Definition 4.11 *An allocation rule A is **monotonic** if for all i, for all $r_i, r_i' \in \mathbb{R}^n$, $r_{-i} \in \mathbb{R}^{(m-1)\times n}$ with $r_{-i} = r'_{-i}$, and their corresponding assignment vectors a_i, a_i' respectively, it holds that*

$$r_i \cdot a_i(r_i, r_{-i}) - r_i' \cdot a_i(r_i, r_{-i}) \;\geqslant\; r_i \cdot a_i(r_i', r_{-i}) - r_i' \cdot a_i(r_i', r_{-i}).$$

Remark 4.12 *The above definition can also be rewritten as follows. An allocation rule A is monotonic if for all i, all $r_i, r_i' \in \mathbb{R}^n$, and all r_{-i}, let o and o' be such that $a_{i,o}(r_i, r_{-i}) = 1$ and $a_{i,o'}(r_i', r_{-i}) = 1$. Then it holds that*

$$r_{i,o} - r_{i,o'} \geqslant r'_{i,o} - r'_{i,o'}.$$

Lemma 4.13 *An allocation rule A is implementable if and only if A is monotonic.*

The characterization follows from specializing the results of Lavi et al. (2003) to our setting.

4.3 A TU-Maximizer Characterization

This section studies the relationship between total utility maximization and the properties introduced above. We begin by showing that TU-maximizing allocation rules satisfy monotonicity, value-creation, and IUO. Then, we establish a structural characterization of the number of matches generated by maximal-cardinality TU-maximizers.

Lemma 4.1 *A TU-maximizer satisfies monotonicity, value-creation and IUO.*

The next theorem establishes our main result.

Theorem 4.2 *Let A be an allocation rule that is:*

1. *monotonic*
2. *value-creating*
3. *irrelevant of unassigned objects, and*
4. *$|A(r)| \geqslant k^{TU}(r)$ for every preference profile r.*

Then, $|A(r)| = k^{TU}(r)$ for all r.

Proof. Suppose, by contradiction, that there exists an allocation rule A which satisfies all of the above conditions, but that for some profile r, we have $|A(r)| > k^{TU}(r)$. We will construct a sequence of modified valuation profiles $r \to r' \to r'' \to r^*$ such that such that $k^{TU}(r^*) > |A(r^*)|$ which contradicts condition 4. We prove this in a sequence of three claims.

Fix an arbitrary matching $A^{TU}(r) \in \mathcal{T}^{TU}(r)$.

Claim 1 *Let $r' = (r'_1, \ldots, r'_m)$ where for all (i, o),*

$$r'_{i,o} = \begin{cases} r_{i,o} & \text{if } (i,o) \in A(r) \cup A^{TU}(r) \\ -1 & \text{otherwise.} \end{cases}$$

Then $A(r') = A(r)$, and $A^{TU}(r') = A^{TU}(r) \in \mathcal{T}^{TU}(r')$.

Proof of Claim 1. Consider r' as defined. Since A and A^{TU} are value-creating, if $(i,o) \notin A(r) \cup A^{TU}(r)$, then $r'_{i,o} = -1 < 0$. Since both A and A^{TU} are value-creating, object o is not assigned to agent i under either allocation rule. Hence, by IUO and Remark 4.4, replacing $r_{i,o}$ by -1 does not affect either allocation. Applying this argument successively to every pair (i,o) such that $(i,o) \notin A(r) \cup A^{TU}(r)$ yields

$$A(r') = A(r) \text{ and } A^{TU}(r') = A^{TU}(r) \in \mathcal{T}^{TU}(r').$$

End proof of Claim 1.

From now on, we work with r as defined in Claim 1 and assume that $r_{i,o} = -1$ for all $(i,o) \notin A(r) \cup A^{TU}(r)$. Given r, create the directed bipartite graph $\mathcal{G}(r)$ as follows.

The set of vertices is $V = M \cup N$. The edge set E is defined by:

- For each $(i,o) \in A(r)$, $(i,o) \in E$.
- For each $(i,o) \in A^{TU}(r)$, $(o,i) \in E$.

Figure 4.1 illustrates an example of the directed bipartite graph $\mathcal{G}(r)$ constructed above.

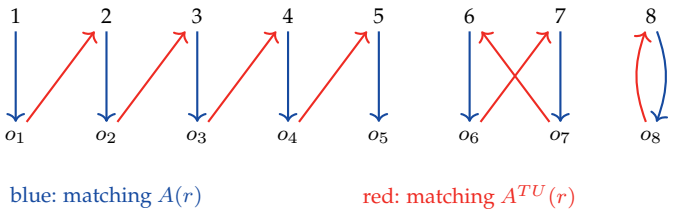

Figure 4.1: Example of the directed bipartite graph $\mathcal{G}(r)$.

Claim 2 *There exists a profile r'' such that,*

1. *there is no maximal path in $\mathcal{G}(r'')$ that starts at an agent and ends at an object, in particular, by Lemma 4.8, $|A(r'')| = k^{TU}(r'')$, and*
2. *$A(r'') \notin \mathcal{T}^{TU}(r'')$.*

Proof of Claim 2. We construct r'' in two steps, corresponding to Claims 2.1 and 2.2. Notice that $\mathcal{G}(r)$ is a directed bipartite graph of the same form as in Lemma 4.8. Hence, by Lemma 4.8, there exists at least one maximal path that starts at an agent and ends at an object in $\mathcal{G}(r)$. Take such a maximal path. We index the agents along this path consecutively, starting from 1, and denote the resulting maximal path as

$$\mathcal{P} = ((1, o_1), (o_1, 2), \ldots, (L, o_L)).$$

Notice that for $1 \leqslant j \leqslant L-1$, object o_j is allocated to agent j under $A(r)$ and to agent $j+1$ under $A^{TU}(r)$. Furthermore, we know from Lemma A.2 that along the maximal path,

$$\sum_{j=1}^{k} r_{j+1,o_j} \;\geqslant\; \sum_{j=1}^{k} r_{j,o_j}, \quad \text{for all } 1 \leqslant k \leqslant L-1. \tag{4.1}$$

Claim 2.1 *There exists a preference profile $\widehat{r}$ such that*

1. *$A(\widehat{r}) = A(r)$,*
2. *$k^{TU}(\widehat{r}) = k^{TU}(r)$, and*
3. *in $\mathcal{G}(\widehat{r})$, all inequalities in (4.1) are strict for all maximal paths that start at an agent and end at an object.*

Proof of Claim 2.1. Suppose that equality in (4.1) holds for at least one $1 \leqslant k \leqslant L-1$. We start by constructing $\widehat{r}$. Choose the biggest K such that

$$\sum_{j=1}^{K} r_{j+1,o_j} \;=\; \sum_{j=1}^{K} r_{j,o_j}.$$

First, we argue that $K \leqslant L-2$. By contradiction, assume that $K = L-1$. Then,

$$\begin{aligned} &\sum_{j=1}^{L-1} r_{j+1,o_j} \;=\; \sum_{j=1}^{L-1} r_{j,o_j} \\ \iff &\sum_{j=1}^{L-1} r_{j+1,o_j} < \sum_{j=1}^{L-1} r_{j,o_j} + r_{L,o_L}. \end{aligned}$$

This contradicts the optimality of a TU-maximizer.

Then, consider agent $K+1$. Since $K \leqslant L-2$, agent $K+1$ exists. Define agent $K+1$'s new preference $\widehat{r}_{K+1} = (\widehat{r}_{K+1,1}, \ldots, \widehat{r}_{K+1,n})$ by

$$\widehat{r}_{K+1,o} = \begin{cases} -1 & \text{if } o = o_K \\ r_{K+1,o} & \text{if } o \neq o_K. \end{cases}$$

Write $\widehat{r} = (\widehat{r}_{K+1}, r_{-(K+1)})$. We then argue that $A(\widehat{r}) = A(r)$. Since agent $K+1$ does not get object o_K under A, by IUO, $A(\widehat{r}) = A(r)$.

Write $T = \{(i, o_i) \mid 1 \leqslant i \leqslant K\} \cup \{(j-1, o_j) \mid K+2 \leqslant j \leqslant L\}$ and denote the set of all agents along the maximal path $\mathcal{P}$ as $M_{\mathcal{P}}$. Define $\widehat{T} = T \cup \{(i, o) \in A^{TU}(r) \mid i \notin M_{\mathcal{P}}\}$. Notice that $K \leqslant L-2$, so $\{(j-1, o_j) \mid K+2 \leqslant j \leqslant L\} \neq \varnothing$. We prove that $\widehat{T} \in \mathcal{T}^{TU}(\widehat{r})$.

A. We show that $\widehat{T}$ maximizes total utility at $\widehat{r}$. By the choice of K, we know that $\sum_{j=1}^{K} r_{j+1,o_j} = \sum_{j=1}^{K} r_{j,o_j}$. Hence, $U(\widehat{T}, \widehat{r}) = U(A^{TU}(r), r)$. Now take any matching T'. We know that $U(A^{TU}(r), r) \geqslant U(T', r)$. Since only r_{K+1,o_K} has changed and decreased, the total utility of any matching that does not assign $(K+1, o_K)$ is unchanged. For matchings that assign $(K+1, o_K)$, total utility decreases compared to the original profile r. Therefore, $U(T', r) \geqslant U(T', \widehat{r})$. Hence, $U(\widehat{T}, \widehat{r}) \geqslant U(T', \widehat{r})$ for every matching T'. Therefore, $\widehat{T}$ maximizes total utility at $\widehat{r}$.

B. Then, we show that $|\widehat{T}| = k^{TU}(\widehat{r})$. Suppose, by contradiction, that there exists a matching T' such that $U(T', \widehat{r}) = U^{TU}(\widehat{r})$ and $|T'| > |\widehat{T}|$. Since only r_{K+1,o_K} changed and decreased, we have $U(T', r) \geqslant U(T', \widehat{r}) = U^{TU}(\widehat{r}) = U(\widehat{T}, \widehat{r}) = U(A^{TU}(r), r)$. Hence, T' also maximizes total utility at r. But then $|T'| > |\widehat{T}| = k^{TU}(r)$, contradicting the definition of $k^{TU}(r)$. Therefore, $k^{TU}(\widehat{r}) = k^{TU}(r)$, and $\widehat{T} \in \mathcal{T}^{TU}(\widehat{r})$.

We transform the original maximal path into a collection of 2-cycles and a new shorter maximal path with no equality. This step is applied inductively to all maximal paths that start at an agent and end at an object where there exists at least one equality along the path. In the end, we obtain a profile $\widehat{r}$ where in $\mathcal{G}(\widehat{r})$, all inequalities are strict. End proof of Claim 2.1.

Figure 4.4 illustrates the above transformation, where a maximal path $\mathcal{P}$ in $\mathcal{G}(r)$ (Figure 4.4(a)) is transformed into 2-cycles and a shorter maximal path (with no equality) in $\mathcal{G}(\widehat{r})$ (Figure 4.4(b)).

Claim 2.2 *Let r be as in Claim 2.1. There exists a preference profile r'' such that $|A(r)| \geqslant |A(r'')| \geqslant k^{TU}(r'') = k^{TU}(r) + 1$.*

Proof of Claim 2.2. Before constructing profile r'', we need an intermediate step.

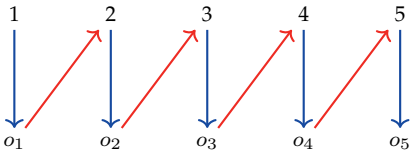

Figure 4.2: A maximal path $\mathcal{P}$ (e.g., from 1 to o_5).

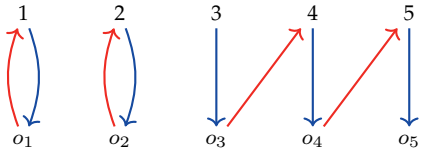

Figure 4.3: After modification: 2-cycles $(1, o_1)$, $(2, o_2)$, and a shorter maximal path from 3 to o_5 (here $K = 2$).

Figure 4.4: Transformation of a maximal path $\mathcal{P}$ in $\mathcal{G}(r)$ into 2-cycles and a shorter maximal path in $\mathcal{G}(\widehat{r})$ in Claim 2.1 (here $K = 2$).

A. Consider agent 1. Write

$$\varepsilon = \min_{1 \leqslant k \leqslant L-1} \left\{ \left(\sum_{j=1}^{k} r_{j+1,o_j} \right) - \left(\sum_{j=1}^{k} r_{j,o_j} \right), \left(\sum_{j=1}^{L-1} r_{j+1,o_j} \right) - \left(\sum_{j=1}^{L} r_{j,o_j} \right) \right\}.$$

Since all inequalities are strict and by Lemma A.1, $\varepsilon > 0$. Define $r'_1 = (r'_{1,1}, \ldots, r'_{1,n})$ by

$$r'_{1,o} = \begin{cases} r_{1,o_1} + \frac{\varepsilon}{8} & \text{if } o = o_1 \\ -\frac{\varepsilon}{8} & \text{if } o = o_L \\ r_{1,o} & \text{if } o \notin \{o_1, o_L\}. \end{cases}$$

Write $r' = (r'_1, r_{-1})$. To prove that $A(r') \notin \mathcal{T}^{TU}(r')$, it suffices to show that $A(r')$ does not maximize total utility at r'. We do this in two steps.

A1. We first show that $(1, o_1) \in A(r')$. Due to value-creation and the fact that $r_{1,o} < 0$ for all $o \neq o_1$, $A(r')$ cannot allocate any object other than o_1 to agent 1. Since $r'_{1,o_1} = r_{1,o_1} + \frac{\varepsilon}{8} > r_{1,o_1}$, due to monotonicity, $A(r')$ must allocate o_1 to agent 1. Hence

$(1, o_1) \in A(r')$.

A2. We now show that $A(r')$ does not maximize total utility at r'. First, we argue that agent 1 cannot remain unmatched while some other matches change. Otherwise, any such change could be replicated in the original TU-maximizing allocation to yield strictly higher total utility, contradicting optimality. Hence, this situation cannot occur.

We know that agent 1 can only get object o_1. Recall from Claim 1 that agent 2 either is not matched or can only get object o_2. Consider the first case. By construction of ε, we know that $r_{2,o_1} > r_{1,o_1} + \frac{\varepsilon}{8}$. Hence, if agent 2 is not matched, giving object o_1 to agent 2 instead of agent 1 yields strictly higher total utility. Therefore, agent 2 must be matched and get object o_2 to increase total utility. Then, agent 3 either is not matched or gets object o_3. Consider the first case. By construction of ε, we know that $r_{2,o_1} + r_{3,o_2} > r_{1,o_1} + \frac{\varepsilon}{8} + r_{2,o_2}$. Hence, if agent 3 is not matched, giving object o_1 to agent 2 and o_2 to agent 3 yields strictly higher total utility than giving object o_1 to agent 1 and o_2 to agent 2. So, agent 3 must get object o_3. By induction, if agent 1 gets object o_1, then agent i must get object o_i for all $2 \leqslant i \leqslant L$. However, by construction of ε,

$$\sum_{j=1}^{L-1} r_{j+1,o_j} > \sum_{j=1}^{L} r_{j,o_j} + \frac{\varepsilon}{8}.$$

Therefore, the matching assigning object o_j to agent $j+1$ for all $1 \leqslant j \leqslant L-1$ yields strictly higher total utility than $A(r')$. Hence, $U(A(r'), r') < U^{TU}(r')$, so $A(r') \notin \mathcal{T}^{TU}(r')$. On the other hand, the matching $A^{TU}(r)$ remains a maximal-cardinality TU-maximizer at r'. Hence, $A^{TU}(r) \in \mathcal{T}^{TU}(r')$.

B. Now, we construct the preference profile r''. Consider again agent 1. Define $r''_1 = (r''_{1,1}, \ldots, r''_{1,n})$ by

$$r''_{1,o} = \begin{cases} r'_{1,o_1} + \frac{3\varepsilon}{8} = r_{1,o_1} + \frac{\varepsilon}{2} & \text{if } o = o_1 \\ \frac{\varepsilon}{8} & \text{if } o = o_L \\ r'_{1,o} & \text{if } o \notin \{o_1, o_L\}. \end{cases}$$

Write $r'' = (r''_1, r_{-1})$.

B1. We show that $|A(r'')| \leqslant |A(r)|$. Since $r''_{1,o_1} - r'_{1,o_1} = \frac{3\varepsilon}{8} > \frac{\varepsilon}{4} = r''_{1,o_L} - r'_{1,o_L}$, monotonicity requires that object o_1 is still allocated to agent 1 under $A(r'')$. Hence, $(1, o_1) \in A(r'')$. Furthermore, by Claim 1 and the value-creation property of the allocation rule A, it follows that $|A(r'')| \leqslant |A(r)|$. To see this, note that in Claim 1, we assign a value of -1 to all agent–object pairs that are not matched under either A^{TU} or A at r. Since A is value-creating, it never assigns pairs with negative value. Hence, under the modified instance r'', A cannot create any new matches involving

these pairs. As a result, there is no way for A to increase the number of matches relative to r, implying that $|A(r'')| \leqslant |A(r)|$.

B2. Let $\tilde{T} := A^{TU}(r) \cup \{(1, o_L)\}$. We now show that $\tilde{T} \in \mathcal{T}^{TU}(r'')$.

Under r'', the only objects that can be assigned to agent 1 without violating value-creation are o_1 and o_L. Consider the case where agent 1 gets o_1. Then, due to Claim 1, agent 2 either is not matched or gets object o_2. Consider the first case. By construction of ε, we know that $r_{2,o_1} > r_{1,o_1} + \frac{\varepsilon}{2}$. Hence, if agent 2 is not matched, giving object o_1 to agent 2 instead of agent 1 yields strictly higher total utility. Therefore, agent 2 must be matched and get object o_2 to maximize total utility. Then, agent 3 either is not matched or gets object o_3. Consider the first case. By construction of ε, we know that $r_{2,o_1} + r_{3,o_2} > r_{1,o_1} + \frac{\varepsilon}{2} + r_{2,o_2}$. Hence, if agent 3 is not matched, giving object o_1 to agent 2 and o_2 to agent 3 yields strictly higher total utility than giving object o_1 to agent 1 and o_2 to agent 2. So, agent 3 must get object o_3. By induction, if agent 1 gets object o_1, then agent i must get object o_i for all $2 \leqslant i \leqslant L$. However, by construction of ε,

$$\begin{aligned} &\sum_{j=1}^{L-1} r_{j+1,o_j} \;>\; \sum_{j=1}^{L} r_{j,o_j} + \frac{\varepsilon}{2} \\ \implies &\sum_{j=1}^{L-1} r_{j+1,o_j} + \frac{\varepsilon}{8} \;>\; \sum_{j=1}^{L} r_{j,o_j} + \frac{\varepsilon}{2}. \end{aligned}$$

Hence, there exists a matching $\widetilde{T}$ obtained from $A^{TU}(r)$ and assigning object o_L to agent 1 such that $U(\widetilde{T}, r'') > U(A(r''), r'')$. Moreover, $|\widetilde{T}| = |A^{TU}(r)|+1 = k^{TU}(r)+1$. Therefore, $k^{TU}(r'') \geqslant k^{TU}(r) + 1$. End proof of Claim 2.2.

Figure 4.7 illustrates how a maximal path $\mathcal{P}$ in $\mathcal{G}(r')$ is transformed into a cycle by adding the edge $(1, o_5)$ in $\mathcal{G}(r'')$.

In the end, we obtain $|A(r'')| \geqslant k^{TU}(r'')$. If the inequality is strict, we know from Lemma 4.8 that there exists at least one maximal path that starts at an agent and ends at an object. We repeat this procedure for all maximal paths that start at an agent and end at an object. This process eventually yields $|A(r'')| = k^{TU}(r'')$, completing the proof of Claim 2.

Claim 3 *There exists a profile r^* such that $|A(r^*)| = |A(r'')|-1$, and $k^{TU}(r^*) = k^{TU}(r'')$. In particular, $|A(r^*)| < k^{TU}(r^*)$.*

Proof of Claim 3. Take r'' as in Claim 2. By construction, after repeatedly applying the transformation described in Claim 2.1, the final maximal path $\mathcal{P} = ((1, o_1), (o_1, 2), \ldots, (L, o_L))$ is transformed into a cycle by assigning object o_L to agent 1 under $A^{TU}(r'')$. Hence, in $\mathcal{G}(r'')$, we obtain the cycle $\mathcal{C} = ((1, o_1), (o_1, 2), \ldots, (L, o_L), (o_L, 1))$. Consider such a cycle $\mathcal{C}$. We start by constructing r^*.

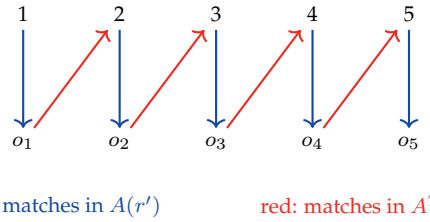

Figure 4.5: A maximal path $\mathcal{P}$ (from 1 to o_5) in $\mathcal{G}(r')$.

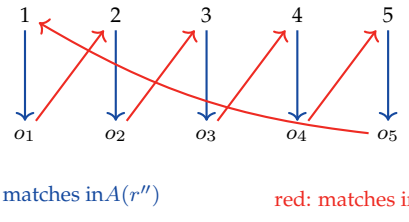

Figure 4.6: Cycle $\mathcal{C}$ in $\mathcal{G}(r'')$ transformed from maximal path $\mathcal{P}$.

Figure 4.7: Transformation of a maximal path $\mathcal{P}$ in $\mathcal{G}(r)$ into a cycle in $\mathcal{G}(r'')$ in Claim 2.2.

Since we only modified the valuation profile of agent 1, according to Lemma A.2, we know that

$$\sum_{j=0}^{k} r_{L-j,o_{L-j-1}} \geqslant \sum_{j=0}^{k} r_{L-j,o_{L-j}}, \text{ for all } 0 \leqslant k \leqslant L-2. \tag{4.2}$$

Choose the biggest K such that $\sum_{j=0}^{K} r_{L-j,o_{L-j-1}} = \sum_{j=0}^{K} r_{L-j,o_{L-j}}$. If all inequalities in (4.2) are strict, write $K = -1$. By the choice of K, we know that $\sum_{j=K+1}^{L-2} r_{L-j,o_{L-j-1}} > \sum_{j=K+1}^{L-2} r_{L-j,o_{L-j}}$.

Write

$$\alpha = \min_{K+1 \leqslant k \leqslant L-3} \left\{ \left(\sum_{j=0}^{k} r_{L-j,o_{L-j-1}} - \sum_{j=0}^{k} r_{L-j,o_{L-j}} \right), \left(\sum_{j=1}^{L-1} r_{j+1,o_j} \right) - \left(\sum_{j=1}^{L} r_{j,o_j} + \frac{\varepsilon}{2} \right) \right\}.$$

Consider agent $L-K-1$. Define a new preference $r^*_{L-K-1} = (r^*_{L-K-1,1}, \ldots, r^*_{L-K-1,n})$

by

$$r^*_{L-K-1,o} = \begin{cases} r_{L-K-1,o_{L-K-2}} - r_{L-K-1,o_{L-K-1}} - \frac{\alpha}{2} & \text{if } o = o_{L-K-2} \\ -\frac{\alpha}{4} & \text{if } o = o_{L-K-1} \\ r_{L-K-1,o} & \text{if } o \notin \{o_{L-K-2}, o_{L-K-1}\}. \end{cases}$$

Write $r^* = (r^*_{L-K-1}, r''_{-(L-K-1)})$.

A. We argue that $|A(r^*)| = |A(r'')| - 1$. Recall that $(1, o_1) \in A(r'')$, and the number of matches under A is the same as $k^{TU}(r'')$. Therefore, due to Claim 1, along the cycle $\mathcal{C}$, $A(r'')$ matches each agent j to object o_j. Since

$$\begin{aligned} r^*_{L-K-1,o_{L-K-2}} - r^*_{L-K-1,o_{L-K-1}} &= r_{L-K-1,o_{L-K-2}} - r_{L-K-1,o_{L-K-1}} - \frac{\alpha}{4} \\ &< r_{L-K-1,o_{L-K-2}} - r_{L-K-1,o_{L-K-1}}, \end{aligned}$$

and $r^*_{L-K-1,o_{L-K-2}} = -\frac{\alpha}{4} < 0$, monotonicity and value-creating require $A(r^*)$ to leave agent $L-K-1$ unmatched. Hence, $|A(r^*)| = |A(r'')| - 1$.

B. Now, we prove that $k^{TU}(r^*) = k^{TU}(r'')$ by showing that $\tilde{T} \in \mathcal{T}^{TU}(r^*)$.

B1. First, we argue that there exists a TU-maximizer in which object o_i is allocated to agent $i+1$ for all $L-K-1 \leqslant i \leqslant L-1$. By construction, $\sum_{j=0}^{K} r_{L-j,o_{L-j-1}} = \sum_{j=0}^{K} r_{L-j,o_{L-j}}$. Hence, replacing the assignments $(L-K-1, o_{L-K-1}), \ldots, (L, o_L)$ by the shifted assignments $(L-K, o_{L-K-1}), \ldots, (L, o_{L-1})$ leaves total utility unchanged. Since the original assignment is TU-maximizing, the shifted assignment is TU-maximizing as well.

B2. Next, we show that any TU-maximizer allocates object o_i is allocated to agent $i+1$ for all $L-K-1 \leqslant i \leqslant L-1$ as in B1 must allocate object o_i to agent $i+1$ for all $1 \leqslant i \leqslant L-K-2$. Assume, by contradiction, that object o_{L-K-2} is allocated to agent $L-K-2$ instead of agent $L-K-1$. However, by construction of α, we know that

$$\sum_{j=0}^{K+2} r_{L-j,o_{L-j-1}} > \sum_{j=0}^{K+2} r_{L-j,o_{L-j}} + \frac{\alpha}{2}.$$

Because $\sum_{j=0}^{K} r_{L-j,o_{L-j-1}} = \sum_{j=0}^{K} r_{L-j,o_{L-j}}$, we have

$$\begin{aligned} &\sum_{j=0}^{K+2} r_{L-j,o_{L-j-1}} > \sum_{j=0}^{K+2} r_{L-j,o_{L-j}} + \frac{\alpha}{2} \\ \iff & r_{L-K-2,o_{L-K-3}} + r_{L-K-1,o_{L-K-2}} > r_{L-K-2,o_{L-K-2}} + r_{L-K-1,o_{L-K-1}} + \frac{\alpha}{2} \\ \iff & r_{L-K-2,o_{L-K-3}} + r_{L-K-1,o_{L-K-2}} - r_{L-K-1,o_{L-K-1}} - \frac{\alpha}{2} > r_{L-K-2,o_{L-K-2}}. \end{aligned}$$

Hence, if a TU-maximizer assigns object o_{L-K-2} to agent $L-K-2$, it must also assign object o_{L-K-3} to agent $L-K-3$. Iterating this argument implies that every TU-maximizer with $(L-K-2, o_{L-K-2})$ must assign object o_j to agent j for all $j < L-K-1$. However, we know that

$$\sum_{j=1}^{L-1} r_{j+1,o_j} > \sum_{j=1}^{L} r_{j,o_j} + \frac{\varepsilon}{2}.$$

And, by construction of α

$$\begin{aligned} &\sum_{j=1}^{L-1} r_{j+1,o_j} > \sum_{j=1}^{L} r_{j,o_j} + \frac{\varepsilon}{2} \\ \iff &\sum_{j=1}^{L-1} r_{j+1,o_j} - r_{L-K-1,o_{L-K-1}} - \frac{\alpha}{2} > \sum_{\substack{j=1 \\ j\neq L-K-1}}^{L} r_{j,o_j} + \frac{\varepsilon}{2}. \end{aligned}$$

Furthermore, since agent 1 and object o_L are both unmatched and $r''_{1,o_L} = \frac{\varepsilon}{8} > 0$, allocating object o_L to agent 1 strictly increases total utility.

B1 and **B2** together imply that $\tilde{T} \in \mathcal{T}^{TU}(r^*)$. End proof of Claim 3.

Figure 4.8 illustrates the final contradiction: A TU-maximizer preserves the entire cycle, while the allocation rule A loses one match by leaving agent $L-K-1$ unmatched.

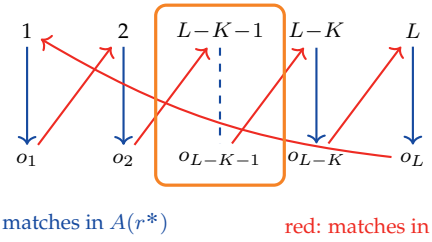

Figure 4.8: The TU-maximizing cycle remains unchanged, while the edge $(L-K-1, o_{L-K-1})$ is removed from $\mathcal{G}(r^*)$, leaving one agent unmatched.

In the end, we find a preference profile r^* such that $|A(r^*)| = |A(r'')|-1 < k^{TU}(r'') = k^{TU}(r^*)$, which contradicts condition 4. ■

4.4 Concluding Remarks

This chapter studied allocation rules for matching markets with heterogeneous cardinal valuations and unit-demand agents. We focused on allocation rules that satisfy

monotonicity, value-creation, and irrelevance of unassigned objects, and compared them to maximal-cardinality total utility maximizers.

Our main result establishes that no allocation rule satisfying these properties can weakly dominate a maximal-cardinality TU-maximizer in terms of the number of matches. In particular, if an allocation rule is value-creating, monotonic, satisfies irrelevance of unassigned objects, and always produces at least as many matches as a maximal-cardinality TU-maximizer, then the two rules must coincide in cardinality at every valuation profile. Thus, under these natural economic and implementability conditions, the matching size generated by TU-maximization is essentially unavoidable. The result highlights the strength of monotonicity restrictions in multidimensional assignment settings. Although implementable rules need not maximize total utility, the combination of monotonicity with value-creation and IUO imposes strong structural constraints on how far such rules can deviate from TU-maximization.

Appendix: Proofs

Proof of Remark 4.4

Since A is value-creating and $v_{i,o} < 0$, object o is not assigned to agent i under v. First suppose that $v'_{i,o} < v_{i,o}$. Then the conclusion follows directly from IUO. Now suppose that $v'_{i,o} > v_{i,o}$. Since $v'_{i,o} < 0$, value-creation also implies that object o is not assigned to agent i under the profile $(v'_{i,o}, v_{-(i,o)})$. Applying IUO to this profile and lowering the report from $v'_{i,o}$ to $v_{i,o}$ yields $A(v'_{i,o}, v_{-(i,o)}) = A(r)$. ■

Proof of Lemma 4.8

Consider $\mathcal{G}(r)$. By construction, each vertex has in-degree and out-degree at most one. Hence, $\mathcal{G}(r)$ decomposes into connected components that are either directed cycles or maximal directed paths.

1. **Cycles.** Every vertex in a cycle has exactly one incoming and one outgoing edge. Hence, cycles contain an equal number of edges from M to N and from N to M.
2. **Maximal paths starting and ending at an object.** These paths have both endpoints in N. They contribute equally to the number of edges from M to N and from N to M.
3. **Maximal paths starting and ending at an agent.** These paths have both endpoints in M. They also contribute equally to both types of edges.
4. **Maximal paths starting at an agent and ending at an object.** These paths begin with an edge from M to N and end with an edge into N. Hence, they contain exactly one more edge from M to N than from N to M.
5. **Maximal paths starting at an object and ending at an agent.** These paths begin with an edge from N to M and end with an edge into M. Hence, they contain exactly one more edge from N to M than from M to N.

Since the total number of edges from M to N is strictly larger than the number of edges from N to M, there must exist at lease one maximal path that starts in M and ends in N. ■

Proof of Lemma 4.1

It is trivial to check value-creation and monotonicity. We now prove that a TU-maximizer satisfies IOU. Let v be a valuation profile and let $A^{TU}(v)$ be a total utility maximizing matching. Fix an agent $i \in M$ and an object $o \in N$ such that $(i, o) \notin A^{TU}(v)$. Consider a modified report v'_i that differs from v_i only in object o, with $v'_{i,o} < v_{i,o}$. All other valuations remain unchanged.

Observe that the total utility of the original matching $A^{TU}(v)$ under the modified profile (v'_i, v_{-i}) remains the same as under v, since (i, o) is not part of the matching. Hence,

$$U(A^{TU}(v), (v'_i, v_{-i})) = U(A^{TU}(v), v).$$

Now take any matching T. We know that $U(A^{TU}(v), v) \geqslant U(T, v)$. Since only $v_{i,o}$ has changed and decreased, the total utility of any matching that does not assign (i, o) is unchanged. For matchings that assign (i, o), total utility strictly decreases relative to the original profile r. Therefore, $U(T, v) \geqslant U(T, (v'_i, v_{-i}))$. It follows that $U(A^{TU}(v), (v'_i, v_{-i})) \geqslant U(T, (v'_i, v_{-i}))$. Hence, $A^{TU}(v)$ remains optimal under (v'_i, v_{-i}), and IUO is satisfied. ■

Lemma A.1 *Let r be such that $|A(r)| > k^{TU}(r)$, and let $\mathcal{G}(r)$ be the graph constructed in the proof of Theorem 4.2 using $A(r)$ and a maximal-cardinality TU-maximizer $A^{TU}(r)$. Suppose that $\mathcal{P}$ is a maximal path in $\mathcal{G}(r)$ that starts at an agent and ends at an object. Then the total utility generated by the matches along $\mathcal{P}$ under $A^{TU}(r)$ is strictly greater than under $A(r)$. That is,*

$$\sum_{j=1}^{L-1} r_{j+1,o_j} \;>\; \sum_{j=1}^{L} r_{j,o_j}.$$

Proof. Assume, by contradiction, that $\sum_{j=1}^{L-1} r_{j+1,o_j} \leqslant \sum_{j=1}^{L} r_{j,o_j}$. Then, replacing the allocations of $A^{TU}(r)$ along $\mathcal{P}$ with those of $A(r)$ would yield a feasible allocation with weakly higher total utility and strictly more matched agents. This contradicts the optimality of $A^{TU}(r)$ as a maximal-cardinality TU-maximizer. ■

Lemma A.2 *Let r be such that $|A(r)| > k^{TU}(r)$, and let $\mathcal{G}(r)$ be the graph constructed in the proof of Theorem 4.2 using $A(r)$ and a maximal-cardinality TU-maximizer $A^{TU}(r)$. Suppose that $\mathcal{P}$ is a maximal path in $\mathcal{G}(r)$ that starts at an agent and ends at an object. Then, along the maximal path $\mathcal{P}$,*

$$\sum_{j=1}^{k} r_{j+1,o_j} \;\geqslant\; \sum_{j=1}^{k} r_{j,o_j}, \quad \textit{for all } 1 \leqslant k \leqslant L-1, \tag{4.1}$$

and

$$\sum_{j=0}^{k} r_{L-j,o_{L-j-1}} \;\geqslant\; \sum_{j=0}^{k} r_{L-j,o_{L-j}}, \quad \textit{for all } 0 \leqslant k \leqslant L-2. \tag{4.2}$$

Proof. We prove the first inequality. The second follows by a symmetric argument applied to the reversed path. Fix $k \in \{1, \ldots, L-1\}$. Assume, by contradiction, that $\sum_{j=1}^{k} r_{j+1,o_j} < \sum_{j=1}^{k} r_{j,o_j}$.

Consider the allocation obtained from an allocation $A^{TU}(r) \in \mathcal{T}^{TU}(r)$ by reassigning the objects $o_1, \ldots, o_k$ to agents $1, \ldots, k$, respectively, while keeping all other assignments unchanged. This reassignment is feasible, since along the maximal path each object o_j is assigned to agent $j+1$ under $A^{TU}(r)$ and to agent j under $A(r)$, and no other agents are affected. Under this modified allocation, the total utility on these k pairs increases strictly by assumption. Hence, the total utility of the new allocation

strictly exceeds that of $A^{TU}(r)$. This contradicts the optimality of $A^{TU}(r)$ as a TU-maximizer. ■

5

Serial Top Trading Cycles and the Exclusion Core in Mixed-Ownership Object Allocation Problems

This chapter is based on joint work with Bettina Klaus and Yang Chen (University of Lausanne).

Earlier versions of this work were presented at the Lausanne Matching and Market Design Workshop 2025, 18th Meeting of the Society for Social Choice and Welfare (SSCW 2026), European Meeting on Game Theory 2026 (SING21), 8th International Workshop on Matching Under Preferences (MATCH-UP 2026), and the MLSE at Maastricht University.

5.1 Introduction

Object allocation problems with indivisible goods arise in many economic settings, including housing markets, school choice, and the assignment of positions or resources when monetary transfers are infeasible or undesirable. Two benchmark models structure much of this literature. In **Shapley–Scarf housing problems** (L. Shapley & Scarf, 1974), each agent initially owns exactly one object, and ownership is purely private. In **house allocation problems** (Hylland & Zeckhauser, 1979), all objects are publicly owned. **House allocation problems with existing tenants** (Abdulkadiroğlu & Sönmez, 1999) lie between these two models: some agents initially own exactly one object, the remaining objects are publicly owned, and other agents own none. These benchmark models are important, but many economically relevant settings involve more complex endowment structures.

In such environments, an agent may own several objects or none, and several agents may jointly own an object. Despite this flexibility in endowments, each agent ultimately receives at most one object. Balbuzanov and Kotowski (2019) call these problems "**simple economies**"; we call them **object allocation problems with complex endowments**. Ownership may be private, public, or shared among agents.

Standard core concepts are grounded in the interpretation that owners can consume or trade their endowed objects. Beyond the benchmark models, the cores built on this interpretation lose their appeal: the strong core may be empty and the weak core may contain wasteful allocations. Balbuzanov and Kotowski instead interpret ownership as a system of exclusion rights: owners can exclude other agents from an object, either directly or through chains of dependence. Their **exclusion core** restricts blocking to deviations that are feasible through these rights. It is always nonempty and contains only Pareto efficient allocations.

Related recent work has further developed the analysis of object allocation with complex endowments. Balbuzanov and Kotowski (2024) extend the exclusion core to economies with production and generalize the top trading cycles algorithm to that setting. Sun et al. (2026) introduce two extensions of the strong core, the effective core and a refinement of it called the rectified core, and show that the effective core neither contains nor is contained in the exclusion core. Zhang (2026) applies the consistency principle to core concepts in this setting and introduces the refined exclusion core, a consistent refinement that admits unaffected agents into blocking coalitions under restrictions. Ishida and Park (2026) study partitioned ownership and its implications for traditional core concepts and the exclusion core.

We focus on **mixed-ownership object allocation problems**, a subclass of object allocation problems with complex endowments. In this subclass, each object is either individually owned by one agent or publicly owned by all agents. Proper co-ownership is ruled out, but an agent may own several objects. This class includes all three benchmark models described above. Importantly, mixed ownership does not eliminate the

difficulties identified by Balbuzanov and Kotowski (2019): even without proper co-ownership, the strong core may still be empty and the weak core may still contain wasteful allocations (Example 5.8). The exclusion core therefore remains an appropriate stability concept.

To analyze exclusion core stability in mixed-ownership object allocation problems, we introduce the **serial top trading cycles (STTC) algorithm**. Given a priority order, an endowment structure, an initial allocation, and agents' preferences, it constructs a sequence of directed graphs and executes top trading cycles sequentially. Unlike the generalized top trading cycles (GTTC) algorithm of Balbuzanov and Kotowski (2019), it allows for an arbitrary initial allocation and uses a single priority order over agents rather than object-specific priority structures. These features make the STTC algorithm the central tool for both generating exclusion-core-stable allocations and analyzing stability from an arbitrary allocation.

Two preliminary results support the main analysis. First, every allocation produced by the STTC algorithm is exclusion core stable (Proposition 5.5). Second, every exclusion-core-stable allocation can be obtained by running the STTC algorithm from the null allocation under a suitable priority order (Proposition 5.6).

We then establish four main results. First, when a properly co-owned object is present, the STTC allocations may form a proper subset of the exclusion core (Theorem 5.1). It follows that, among domains defined by restrictions on endowment structures, the mixed-ownership domain is the largest on which the STTC algorithms characterize the exclusion core (Corollary 5.2). This maximality result explains our focus on mixed ownership. Second, we provide an exclusion-core-stability test (Theorem 5.3). An allocation is exclusion core stable if and only if it is the outcome of the STTC algorithm initialized at that allocation, independently of the priority order. Third, we show that the exclusion core is a von Neumann–Morgenstern (1944) stable set with respect to exclusion-domination (Theorem 5.6). Fourth, we capture the dynamic aspect of exclusion-blocking: from any allocation, there exists a self-contained exclusion-blocking path to the exclusion core (Theorem 5.11). As a consequence, a random process whose transition probabilities depend only on the current allocation and assign positive probability to every available post-blocking allocation reaches an exclusion-core-stable allocation with probability one (Corollary 5.12).

We next clarify how our results relate to those of Balbuzanov and Kotowski (2019). Remark 5.9 extends Proposition 5.5 to the full class of object allocation problems with complex endowments: every allocation generated by an owner-respecting flexible TTC algorithm is exclusion core stable. For algorithms in which each object's edge is determined by a fixed priority order over agents, this also follows from Corollary 1(a) and Lemma 3(a) of Balbuzanov and Kotowski (2019); flexible TTC algorithms are more general, since an object may point to different remaining owners at different steps. Proposition 4 of Balbuzanov and Kotowski (2019) establishes von Neumann–Morgenstern stability for a subclass of these problems. Our Theorem 5.6 is closely

related, but we state and prove it separately because its constructive proof is intertwined with the path result of Section 5.4.4. Remark 5.8 explains how Proposition 5.5 extends when the common priority order is replaced by arbitrary edge choices.

For mixed-ownership problems, we prove von Neumann–Morgenstern stability jointly with the path-to-exclusion-core-stability result. The constructive proof links the myopic, step-by-step self-contained exclusion-blockings forming the path to the farsighted external-stability blocking by the final STTC allocation.

Our results differ in how far they extend beyond mixed ownership. Exclusion core stability of STTC allocations is the most robust: it holds for the full class of object allocation problems with complex endowments (Remark 5.9). The exclusion core characterization and the fixed-point test are the least robust: a single properly co-owned object is enough for both to fail. Call a flexible TTC algorithm for a complex-endowment problem owner-respecting if an object points to one of its owners whenever an owner remains. The construction proving Theorem 5.1 then yields an exclusion-core-stable allocation that is a fixed point of no owner-respecting flexible TTC algorithm.

Von Neumann–Morgenstern stability fails only in part: internal stability holds for every endowment structure, whereas Example E.9 exhibits a complex-endowment problem whose exclusion core is not externally stable. Both external stability and the existence of blocking paths, however, reach beyond mixed ownership. Appendix E shows that an endowment structure is acyclic in the sense of Balbuzanov and Kotowski (2019) exactly when either no object is properly co-owned or all objects share a single set of owners (Proposition E.2), and that both conclusions continue to hold on this larger class (Corollaries E.7 and E.8): Theorems 5.6 and 5.11 cover the first case and Theorems E.5 and E.6 the second. Acyclicity is not necessary for external stability either: Example E.10 exhibits an endowment structure that is not acyclic and for which the exclusion core is externally stable at every preference profile. Whether a self-contained exclusion-blocking path to the exclusion core always exists under proper co-ownership remains open. Nevertheless, such a path exists in both examples in Appendix E: a path of length two in Example E.9 and a path of length one in Example E.10.

The remainder of the paper is organized as follows. Section 5.2 introduces the model and the relevant core concepts. Section 5.3 defines the STTC algorithm and establishes its basic relationship with the exclusion core. Section 5.4 presents our main results, and Section 5.5 concludes. Appendices A–C contain all proofs. Appendix D provides a detailed example illustrating the STTC algorithm and four stability results: the exclusion-core characterization in Proposition 5.6, the exclusion-core-stability test in Theorem 5.3, the external-stability conclusion of Theorem 5.6, and the path construction in Theorem 5.11. Appendix E studies acyclic endowment structures. It characterizes them, shows that Theorems 5.6 and 5.11 extend from the mixed-ownership structures to this larger class, and closes with two complex-endowment examples:

one whose exclusion core fails to be externally stable and one non-acyclic endowment structure whose exclusion core is externally stable at every preference profile.

5.2 Preliminaries

5.2.1 Model

We consider a finite **set of agents** $\boldsymbol{N = \{1, \ldots, n\}}$ and a finite **set of objects** $\boldsymbol{O = \{o_1, \ldots, o_m\}}$. The numbers of agents and objects may differ. **Ownership is mixed**: each object is either individually owned by exactly one agent or publicly owned by all agents. For each agent $i \in N$, we denote the set of objects he individually owns by ω_i. The list $\boldsymbol{\omega = (\omega_1, \ldots, \omega_n)}$ then describes the **endowment structure**. We denote the **set of individually owned objects** by $\boldsymbol{O_I = \bigcup_{i \in N} \omega_i}$ and the set of **public objects** by $\boldsymbol{O_P = O \backslash O_I}$.

A nonempty subset $S \subseteq N$ is called a **coalition**. For each coalition $S \subsetneq N$, the **set of objects collectively owned by** $\boldsymbol{S}$ is denoted by $\boldsymbol{\omega(S)} = \bigcup_{i \in S} \omega_i \subseteq O_I$. For the grand coalition, we define $\boldsymbol{\omega(N) := O}$. Thus, public ownership is represented collectively rather than through the individual sets ω_i: public objects do not belong to any ω_i, but they belong to the endowment of the grand coalition. To illustrate the endowment structure, consider the following mixed-ownership situation with three agents $N = \{1, 2, 3\}$ and four objects $O = \{a, b, c, d\}$. If agent 1 owns objects a and b, agent 2 owns object c, and agent 3 owns no object, then $\omega_1 = \{a, b\}$, $\omega_2 = \{c\}$, $\omega_3 = \varnothing$, $\omega(\{1, 2\}) = \{a, b, c\}$, $\omega(\{1, 3\}) = \{a, b\}$, $\omega(\{2, 3\}) = \{c\}$, and $\omega(\{1, 2, 3\}) = \{a, b, c, d\}$.

Each agent has an outside option, modeled as the **null object** $\boldsymbol{\varnothing}$,[1] which may literally represent "receiving nothing." The null object has enough capacity for all agents. When we need to distinguish the elements of O from the null object, we call them **real objects**.

Balbuzanov and Kotowski (2019) study a more general class of object allocation problems. They call these problems "simple economies"; we call them **object allocation problems with complex endowments**. In their model, objects may be co-owned by multiple agents.[2] Our mixed-ownership model is the subclass in which each object is either individually owned by exactly one agent or publicly owned by all agents.[3]

[1] Throughout, $\varnothing$ denotes the null object, whereas $\emptyset$ denotes the empty set.

[2] Balbuzanov and Kotowski (2019) consider the following more general model. For each object $o \in O$, $S^o \subseteq N$ denotes the set of owners of object o. If an object $o \in O$ is individually owned by agent $i \in N$, then $S^o = \{i\}$, and if it is publicly owned, then $S^o = N$. Balbuzanov and Kotowski (2019), however, also allow proper co-ownership, that is, ownership of an object by a coalition S^o with $1 < |S^o| < n$, which is not allowed in our model. Because this broader model permits complex endowment structures, we use the term "object allocation problems with complex endowments" for the problems introduced by Balbuzanov and Kotowski (2019).

[3] In one result, Balbuzanov and Kotowski (2019, Lemma 3(b)) consider the same model as ours, referring to it as a "simple economy where every house is either privately or publicly owned."

We assume that each agent $i \in N$ has a preference relation R_i over the set of objects and the null object $O \cup \{\varnothing\}$, which is *complete*, *transitive*, and *antisymmetric*. For an agent $i \in N$ and two objects $a, b \in O \cup \{\varnothing\}$, we write $a \; R_i \; b$ if agent i weakly prefers object a to object b, and $a \; P_i \; b$ if agent i strictly prefers object a to object b. Thus, P_i is the strict part of R_i, and, by antisymmetry, $a \, R_i \, b$ if and only if $a \, P_i \, b$ or $a = b$. Let $R := (R_i)_{i \in N}$ denote a preference profile. A **mixed-ownership object allocation problem** (henceforth, a **problem**) is specified by a quadruple (N, O, ω, R). Throughout, the null object $\varnothing$ is implicitly available in every problem. Whenever (N, O, ω) is fixed, we simply denote a problem by R.

We assume that each agent demands exactly one object in $O \cup \{\varnothing\}$. An **allocation** $\mu : N \to O \cup \{\varnothing\}$ assigns each agent $i \in N$ an **allotment** $\mu(i) \in O \cup \{\varnothing\}$, and assigns each object $o \in O$ to at most one agent. We also say that agent i **occupies** his allotment $\mu(i)$ at μ. We also write an allocation as a list of its allotments, $\mu = (\mu_1, \ldots, \mu_n)$, where the i-th entry is agent i's allotment. We denote the set of allotments of the members of a coalition S at allocation μ by $\mu(S) = \bigcup_{i \in S} \{\mu(i)\}$.

Our mixed-ownership model includes several classical object allocation models as special cases:

- when each agent owns exactly one object and the number of agents equals the number of objects, our model reduces to the **Shapley–Scarf housing market model** (L. Shapley & Scarf, 1974);
- when all objects are public, our model reduces to the **house allocation model** (Hylland & Zeckhauser, 1979);
- when some agents own exactly one object and all remaining objects are public, our model reduces to the **house allocation model with existing tenants** (Abdulkadiroğlu & Sönmez, 1999).

5.2.2 Core stability notions

We begin with the classical notions of weak blocking and strong blocking. They define the strong core and the weak core, respectively. This distinction and its consequences for markets with indivisible goods go back to Roth and Postlewaite (1977). These definitions are grounded in the interpretation that owners can freely **consume** or **trade** their endowed objects. In line with this interpretation, we say that a coalition $\boldsymbol{S}$ **is trade-enforcing at allocation** $\boldsymbol{\sigma}$ if $\sigma(S) \subseteq \omega(S) \cup \{\varnothing\}$.

Definition 5.1 (Strong core stability) *Given a problem R, allocations μ and σ, and a coalition S, we say that* $\boldsymbol{S}$ ***weakly blocks*** $\boldsymbol{\mu}$ ***via*** $\boldsymbol{\sigma}$ *if*

1. *coalition S is trade-enforcing at σ; and*
2. *each agent in S is weakly better off at σ, and at least one agent in S is strictly better off at σ; i.e., for each $i \in S$, $\sigma(i) \; R_i \; \mu(i)$, and, for some $j \in S$, $\sigma(j) \; P_j \; \mu(j)$.*

*The **strong core** for a problem R, denoted by $\mathrm{SC}(R)$, consists of all allocations that cannot be weakly blocked by any coalition. An allocation in the strong core is **strongly core stable**.*

Definition 5.2 (Weak core stability) *Given a problem R, allocations μ and σ, and a coalition S, we say that $\boldsymbol{S}$ **strongly blocks** $\boldsymbol{\mu}$ **via** $\boldsymbol{\sigma}$ if*

1. *coalition S is trade-enforcing at σ; and*
2. *each agent in S is strictly better off at σ, i.e., for each $i \in S$, $\sigma(i)\ P_i\ \mu(i)$.*

*The **weak core** for a problem R, denoted by $\mathrm{WC}(R)$, consists of all allocations that cannot be strongly blocked by any coalition. An allocation in the weak core is **weakly core stable**.*

Strong core allocations also satisfy the following classical efficiency criterion.

Definition 5.3 (Pareto efficiency) *Given a problem R and allocations μ and σ, we say that $\boldsymbol{\sigma}$ **Pareto dominates** $\boldsymbol{\mu}$ if each agent in N is weakly better off at σ and at least one agent in N is strictly better off at σ; i.e., for each $i \in N$, $\sigma(i)\ R_i\ \mu(i)$, and, for some $j \in N$, $\sigma(j)\ P_j\ \mu(j)$. An allocation μ is **Pareto efficient** if no allocation σ Pareto dominates it.*

We now introduce exclusion-blocking, following Balbuzanov and Kotowski (2019). The central idea is to interpret ownership as a system of exclusion rights rather than solely as a right to consume or trade endowed objects. Ownership allows a coalition to evict other agents from objects it owns and therefore controls.

For any coalition S, we call $\boldsymbol{\omega}(\boldsymbol{S})$ its **direct eviction set**.[4] It is the set of objects over which coalition S has direct exclusion rights. By exercising these rights, coalition S can directly evict the current occupants of these objects. The set $\omega(S)$ depends only on the endowment structure; the prevailing allocation determines which agents occupy these objects.

By exercising its direct exclusion rights, a coalition may become able to exercise further exclusion rights. Suppose that coalition S can evict agent j from his allotment at allocation μ. Agent j then depends on continued accommodation by S. In the exclusion-core logic, coalition S can therefore induce agent j to exercise the exclusion rights associated with his endowment. This may allow the coalition to evict further agents and thereby gain access to their exclusion rights. We formalize this chain of dependence recursively.

Fix an allocation μ and a coalition S. For each set of objects $A \subseteq O$, let $\mu^{-1}(A)$ denote the set of agents whose allotment at μ belongs to A, that is, the agents occupying the objects in A: $\mu^{-1}(A) := \{i \in N \mid \mu(i) \in A\}$.

[4] In the mixed-ownership model, $\omega(S)$ is defined in Section 5.2.1. For object allocation problems with complex endowments, Footnote 2 defines S^o as the set of owners of object o and $\omega(S) = \{o \in O \mid S^o \subseteq S\}$.

For each $k \geqslant 0$, the set S_k^μ will consist of the original coalition and all agents reached through at most k rounds of eviction. Initially, $S_0^\mu := S$. If the agents in S_k^μ have been reached, the coalition controls the objects in $\omega(S_k^\mu)$ and can therefore reach the agents currently occupying those objects. Accordingly, define $S_{k+1}^\mu := S_k^\mu \cup \mu^{-1}\big(\omega(S_k^\mu)\big)$.

The sequence $S_0^\mu \subseteq S_1^\mu \subseteq S_2^\mu \subseteq \cdots$ is increasing. Since N is finite, it reaches a fixed point. Let ℓ be the smallest integer such that $S_{\ell+1}^\mu = S_\ell^\mu$. The **extended eviction set of coalition S at allocation μ** is $\Omega(S \mid \mu) := \omega(S_\ell^\mu)$. Thus, $\Omega(S \mid \mu)$ contains all objects from which coalition S can evict agents either directly or through a finite chain of evictions.

Endowments interpreted as exclusion rights give rise to the following counterpart of trade-enforcement. Given allocations μ and σ, coalition ***S* exclusion-enforces σ at μ** if every agent outside S who is worse off at σ occupies at μ an object in the extended eviction set of S; formally, for each $j \notin S$, $\mu(j) \, P_j \, \sigma(j)$ implies $\mu(j) \in \Omega(S \mid \mu)$. Every outsider harmed by the move to σ must therefore be removable through exclusion rights available to the coalition, either directly or through the recursive chain of dependence.

Definition 5.4 (Exclusion core stability) *Given a problem R, allocations μ and σ, and a coalition S, we say that **S exclusion-blocks μ via σ** if*

1. *coalition S exclusion-enforces σ at μ; and*
2. *each agent in S is strictly better off at σ, i.e., for each $i \in S$, $\sigma(i) \, P_i \, \mu(i)$.*

*The **exclusion core** for a problem R, denoted by $\mathrm{EC}(R)$, consists of all allocations that cannot be exclusion-blocked by any coalition. An allocation in the exclusion core is **exclusion core stable**.*

Four properties of extended eviction sets recur throughout the analysis: inclusion of the direct eviction set, coalitional monotonicity, closure under the endowments of agents reached through eviction, and propagation along chains of dependence.

Lemma 5.5 (Properties of extended eviction sets) *Let R be a problem, let μ be an allocation, and let S and S' be coalitions.*

(a) $\omega(S) \subseteq \Omega(S \mid \mu)$.
(b) **Monotonicity.** *If $S \subseteq S'$, then $\Omega(S \mid \mu) \subseteq \Omega(S' \mid \mu)$.*
(c) **Closure.** *For an agent $i \in N$, if $\mu(i) \in \Omega(S \mid \mu)$, then $\omega_i \subseteq \Omega(S \mid \mu)$.*
(d) **Chains of dependence.** *Consider agents $i^1, \ldots, i^q$ and real objects $o^1, \ldots, o^q$ arranged in an ownership–occupancy chain: for each $r \in \{1, \ldots, q\}$, agent i^r owns object o^r (that is, $o^r \in \omega_{i^r}$), and, for each $r \in \{1, \ldots, q-1\}$, agent i^r occupies object o^{r+1} (that is, $\mu(i^r) = o^{r+1}$). We depict such a chain by*

$$o^1 \longrightarrow i^1 \longrightarrow o^2 \longrightarrow i^2 \longrightarrow \cdots \longrightarrow o^q \longrightarrow i^q,$$

where an arrow from an object to an agent records that the agent owns the object, and an arrow from an agent to an object records that the object is his allotment at μ. If coalition S reaches agent i^q, in the sense that $i^q \in S$ or $\omega_{i^q} \subseteq \Omega(S \mid \mu)$, then $o^1 \in \Omega(S \mid \mu)$.

Appendix A contains the proof of Lemma 5.5.

The following example illustrates extended eviction sets and exclusion-blocking.

In the preference tables and displayed allocations in the examples below, boldface marks each agent's allotment at the reference allocation, while underlining marks his allotment at the alternative allocation. An allotment shown in both formats is unchanged.

Example 5.6 (Exclusion-blocking) *Consider a Shapley–Scarf housing market with three agents $N = \{1, 2, 3\}$ and three objects $O = \{a, b, c\}$. Agents 1, 2, and 3 own objects a, b, and c, respectively. Preferences are as follows.*

$$\begin{array}{llllll} R_1: & \underline{a} & \boldsymbol{c} & b & \varnothing \\ R_2: & \boldsymbol{a} & \underline{b} & c & \varnothing \\ R_3: & \boldsymbol{b} & a & \underline{c} & \varnothing \end{array}$$

Consider allocation $\boldsymbol{\mu} = (\boldsymbol{c}, \boldsymbol{a}, \boldsymbol{b})$ and the singleton coalition $S = \{1\}$. Initially, $S_0^\mu = S = \{1\}$ and $\omega(S_0^\mu) = \{a\}$. Since agent 2 receives object a at μ, $S_1^\mu = S_0^\mu \cup \mu^{-1}\big(\omega(S_0^\mu)\big) = \{1, 2\}$. Together, agents 1 and 2 own objects a and b, and therefore $\omega(S_1^\mu) = \{a, b\}$. Since agent 3 receives object b at μ, $S_2^\mu = S_1^\mu \cup \mu^{-1}\big(\omega(S_1^\mu)\big) = \{1, 2, 3\} = N$. Hence, $\omega(S_2^\mu) = O = \{a, b, c\}$ and $S_3^\mu = S_2^\mu$. It follows that $\Omega(\{1\} \mid \mu) = \{a, b, c\}$.

What the recursion traces here is a chain of dependence in the sense of Lemma 5.5(d): agent 1 owns object a, which agent 2 occupies; agent 2 owns object b, which agent 3 occupies; and agent 3 owns object c. In the notation of that lemma, $c \longrightarrow 3 \longrightarrow b \longrightarrow 2 \longrightarrow a \longrightarrow 1$, and coalition $\{1\}$ reaches object c by way of agents 2 and 3. As a result, agent 1 can exclusion-block μ via allocation $\underline{\sigma} = (\underline{a}, \underline{b}, \underline{c})$. ◇

Lemma 5.7 (Pareto efficiency and exclusion-blocking) *Let R be a problem. If σ Pareto dominates μ, then the coalition of agents who are strictly better off at σ than at μ exclusion-blocks μ via σ. Hence, if an allocation is exclusion core stable, then it is Pareto efficient.*

Appendix A contains the proof of Lemma 5.7.

5.2.3 Relationships between core concepts

We now describe the relationships among the traditional core concepts and the exclusion core for mixed-ownership object allocation problems. The strong core is a refinement of the weak core, but it may be empty; the weak core is always nonempty, but it may contain wasteful allocations (see Example 5.8). The exclusion core is always nonempty and consists only of Pareto efficient allocations. Moreover, the exclusion core is a refinement of the weak core (Balbuzanov & Kotowski, 2019, Lemma 1): $\mathrm{EC}(R) \subseteq \mathrm{WC}(R)$. However, neither the strong core nor the exclusion core contains the other. Example 5.8 shows that the exclusion core need not be contained in the strong core, whereas Example 5.9 shows that the strong core need not be contained in the exclusion core.

The two sets may nevertheless coincide: in Example 5.6, a Shapley–Scarf housing market, the strong core and the exclusion core both consist of the single allocation assigning object a to agent 1, object b to agent 2, and object c to agent 3.[5] These relationships are depicted in Figure 5.1.

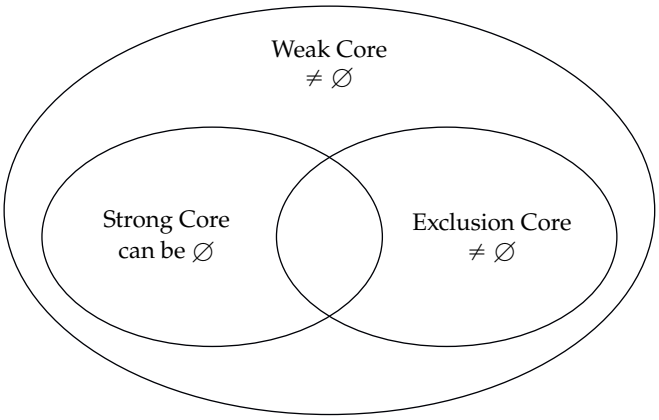

Figure 5.1: The relationships among the weak core, the strong core, and the exclusion core for mixed-ownership object allocation problems.

The following example, from Balbuzanov and Kotowski, 2019, illustrates that the strong core can be empty and that the exclusion core can be a strict subset of the weak core.

Example 5.8 (Empty strong core and strict inclusion of exclusion core) *Consider a problem with three agents $N = \{1, 2, 3\}$ and two objects $O = \{a, b\}$. Assume that agent 1 owns both objects a and b. Preferences are as follows.*

[5] In a housing market with strict preferences, the strong core consists of the unique top trading cycles allocation (Roth & Postlewaite, 1977), which is the allocation just described. That the exclusion core is the same singleton follows from Proposition 5.6.

$$\begin{array}{lcccc} R_1: & \boldsymbol{a} & b & \varnothing & \\ R_2: & a & b & \boldsymbol{\varnothing} & \\ R_3: & a & b & \boldsymbol{\varnothing} & \end{array}$$

Consider allocation $\boldsymbol{\mu} = (\boldsymbol{a}, \boldsymbol{\varnothing}, \boldsymbol{\varnothing})$. *Any coalition that trade-enforces a different allocation must include agent* 1. *Since agent* 1 *receives his most-preferred object at* μ, *no coalition can strongly block* μ. *Thus,* μ *is a weak core allocation. However, since object* b *is not assigned,* μ *is wasteful.*

The strong core in this market is empty. At any strong core allocation, agent 1 *must receive object* a; *otherwise, the singleton coalition* $\{1\}$ *weakly blocks by assigning object* a *to agent* 1. *Pareto efficiency then requires object* b *to be assigned to some agent* $i \in \{2, 3\}$. *Let* j *be the other agent in* $\{2, 3\}$. *Coalition* $\{1, j\}$ *can trade-enforce an allocation at which agent* 1 *keeps object* a *and agent* j *receives object* b. *Agent* 1 *is indifferent and agent* j *is strictly better off, so this coalition weakly blocks the original allocation.*

The exclusion core in this market consists of two allocations μ^1 *and* μ^2 *where agent* 1 *receives object* a, *and object* b *is assigned to either* 2 *or* 3: $\mu^1 = (a, b, \varnothing)$ *and* $\mu^2 = (a, \varnothing, b)$. *Both* μ^1 *and* μ^2 *are weak core allocations, while the allocation* $\mu = (a, \varnothing, \varnothing)$ *considered above is a weak core allocation that does not belong to the exclusion core. Hence, the exclusion core is a strict subset of the weak core.* ◇

The following example, also from Balbuzanov and Kotowski, 2019, shows that the strong core need not be contained in the exclusion core.

Example 5.9 (A strong core allocation that is not exclusion core stable) *Consider a problem with four agents* $N = \{1, 2, 3, 4\}$ *and four objects* $O = \{a, b, c, d\}$. *Agents* 1, 2, *and* 3 *own objects* a, b, *and* c, *respectively;* d *is publicly owned. Preferences are as follows.*

$$\begin{array}{lccccc} R_1: & \underline{b} & \boldsymbol{a} & \varnothing & \ldots & \\ R_2: & \underline{d} & \boldsymbol{c} & b & \varnothing & a \\ R_3: & \boldsymbol{b} & \underline{c} & \varnothing & \ldots & \\ R_4: & \underline{a} & \boldsymbol{d} & c & \varnothing & b. \end{array}$$

Consider allocation $\boldsymbol{\mu} = (\boldsymbol{a}, \boldsymbol{c}, \boldsymbol{b}, \boldsymbol{d})$. *Define* $\underline{\sigma} := (\underline{b}, \underline{d}, \underline{c}, \underline{a})$. *The allocation* μ *belongs to the strong core and hence to the weak core, but not to the exclusion core: coalition* $S = \{1, 2, 4\}$ *exclusion-enforces* σ *at* μ *and therefore exclusion-blocks* μ *via* σ.[6] ◇

[6] The only worse-off agent is agent 3 with $\mu(3) = b \in \Omega(\{1, 2, 4\} \mid \mu) = \{a, b, c, d\}$: the recursion first reaches agents 1 and 3, who occupy the objects in $\omega(\{1, 2, 4\}) = \{a, b\}$, so its terminal agent set is the grand coalition N, and $\omega(N) = O$.

5.3 Serial Top Trading Cycles

5.3.1 STTC graphs and the STTC algorithm

In this subsection, we introduce the STTC graphs and reduced STTC graphs used to define the STTC algorithm. For a given problem, priority order, and arbitrary initial allocation, the STTC algorithm serves as the central tool in our analysis of exclusion core stability.

Let $\boldsymbol{\triangleright}$ denote a **priority order** (i.e., a strict order) over the set of agents N, and let $\boldsymbol{\Delta}$ denote the set of priority orders over N. Write $i \triangleright j$ when agent i has higher priority than agent j.

Given a problem (N, O, ω, R), an allocation μ, and a priority order $\triangleright$ over the set of agents N, we refer to $(\boldsymbol{N, O, \omega, R, \mu, \triangleright})$ as an **instance**.

Definition 5.1 (Serial top trading cycles graphs) *Given an instance $(N, O, \omega, R, \mu, \triangleright)$, let $\boldsymbol{V} = (N \cup O \cup \{\varnothing\})$ be the **set of vertices** and $\boldsymbol{E}$ the **set of edges**. We denote a directed edge from vertex x to vertex y by (x, y) and say that x points to y. The set of edges E is defined as follows:*

- *Each agent $i \in N$ points to his most-preferred object o in $O \cup \{\varnothing\}$ according to R_i, yielding an edge (i, o);*
- *each individually owned object $o \in O_I$ points to its owner i, yielding an edge (o, i);*
- *each public object $o \in O_P$ that is assigned at μ points to its assigned agent i at μ (i.e., $\mu(i) = o$), yielding an edge (o, i);*
- *each public object $o \in O_P$ that is unassigned at μ points to the agent i with the highest priority according to $\triangleright$, yielding an edge (o, i);*
- *the null object $\varnothing$ points to each agent $i \in N$ who points to $\varnothing$, yielding an edge $(\varnothing, i)$ for each such agent.*

*We call the resulting directed graph the **serial top trading cycles (STTC) graph***

$$\boldsymbol{\mathcal{G}(N, O, \omega, R, \mu, \triangleright)} := (V, E)$$

(or $\mathcal{G}(R, \mu, \triangleright)$ for short).

For an individually owned object, the edge (o, i) records exactly the ownership relation appearing in the chains of dependence of Lemma 5.5(d), which is what allows Lemma B.1 in Appendix B to read such a chain directly off a cycle of the algorithm.

A **cycle $\boldsymbol{\mathcal{C}}$** in an STTC graph, such as $\mathcal{G}(R, \mu, \triangleright)$, is an alternating sequence of agents and objects connected by directed edges: $o^1 \to i^1 \to o^2 \to \cdots \to o^k \to i^k \to o^1$. Denote the set of all agents who participate in such a cycle $\mathcal{C}$ by $S(\mathcal{C})$. A cycle containing a single agent, $i \to o \to i$, is a **self-cycle**; its object o may be a real object or the null

object. Independently of this, we call a cycle **ordinary** if it does not involve the null object, that is, if all of its objects are real.

Since the null object points to *every* agent who points to it, a sequence through $\varnothing$ could formally close through several agents; we therefore adopt the convention that each agent pointing to $\varnothing$ forms the self-cycle $i \to \varnothing \to i$, and that this is the only cycle through the null object containing agent i. Every vertex of an ordinary cycle has out-degree one. Hence, distinct cycles are always disjoint in agents and in real objects; the only vertex that distinct cycles can share is the null object, which has enough capacity for all agents.

The STTC algorithm proceeds in steps. At each step, all agents belonging to cycles leave with the objects they receive, and the same graph construction is then applied to the remaining agents and objects. The following definition formalizes the resulting reduced instance. Its components restrict the original data to the surviving agents and objects, with one substantive exception: an agent whose allotment has already been taken is recorded at the reduced allocation as holding the null object.

Definition 5.2 (Reduced instances and STTC graphs) *Consider $(N, O, \omega, R, \mu, \rhd)$. Suppose that the agents in a set $\bar{N}$ leave the problem with their allotments under an allocation σ, which may differ from μ. The remaining sets of agents and objects are $\boldsymbol{N'} = N \backslash \bar{N}$ and $\boldsymbol{O'} = O \backslash \sigma(\bar{N})$; as always, the null object remains available.*

- *If $N' = \{i\}$, then define $\omega'_i = O'$; otherwise, for each $i \in N'$, let $\omega'_i := \omega_i \cap O'$.*[7] *Then, the **reduced endowment structure** is $\boldsymbol{\omega'} = (\omega'_i)_{i \in N'}$. Note that $\bigcup_{i \in N'} \omega'_i \subseteq O'$.*
- *For each agent $i \in N'$, let R'_i be the restriction of R_i to $O' \cup \{\varnothing\}$; that is, for all $a, b \in O' \cup \{\varnothing\}$, $a\ R'_i\ b$ if and only if $a\ R_i\ b$. Then, the **reduced preference profile** is $R' := (R'_i)_{i \in N'}$.*
- *For each $i \in N'$, let*

$$\mu'(i) := \begin{cases} \mu(i) & \text{if } \mu(i) \in O' \cup \{\varnothing\} \\ \varnothing & \text{otherwise.} \end{cases}$$

  *The resulting **reduced allocation** is $\boldsymbol{\mu'} = (\mu'(i))_{i \in N'}$.*
- *For each pair $i, j \in N'$, let $i \rhd' j$ if and only if $i \rhd j$. Then, $\boldsymbol{\rhd'}$ denotes the **reduced priority order** over the set of agents N'.*

*The resulting **reduced instance** is $(\boldsymbol{N', O', \omega', R', \mu', \rhd'})$, and its **reduced STTC graph** is $\boldsymbol{\mathcal{G}(N', O', \omega', R', \mu', \rhd')}$.*

We now use STTC graphs and their reductions to define the STTC algorithm.

[7] When only one agent remains, private and public ownership induce the same edges in the reduced STTC graph: every remaining object points to the remaining agent, either as its owner or, if it is public, as its assignee or as the highest-priority remaining agent. Thus, setting $\omega'_i = O'$ is without loss and simplifies the definition of the reduced instance. Note, however, that $\omega'_i \subseteq \omega_i$ holds whenever $|N'| \geqslant 2$, but fails when $N' = \{i\}$.

Definition 5.3 (The serial top trading cycles algorithm)

Input: *An instance* $(N, O, \omega, R, \mu, \rhd)$.

Step 1*. Let* $N^1 := N$ *and* $O^1 := O$*, and construct the STTC graph* $\mathcal{G}^1 = \mathcal{G}(N^1, O^1, \omega, R, \mu, \rhd)$ *according to Definition 5.1.*

Identify all top trading cycles in $\mathcal{G}^1$ *and execute each cycle:*[8]

- *For each cycle* $\mathcal{C}$*, let* $S(\mathcal{C})$ *denote its set of agents. Assign each* $i \in S(\mathcal{C})$ *the object to which he points in* $\mathcal{C}$*, and denote this allotment by* $\mathrm{STTC}_i(R, \mu, \rhd)$.
- *Let* $\bar{N}^1$ *and* $\bar{O}^1$ *denote the sets of agents and objects, respectively, belonging to cycles. Remove these agents and objects, except for the null object* $\varnothing$*, from the current instance, and set* $N^2 := N^1 \backslash \bar{N}^1$ *and* $O^2 := O^1 \backslash \bar{O}^1$.
- *If* $N^2 = \varnothing$*, stop. If* $N^2 \neq \varnothing$ *and* $O^2 = \varnothing$*, assign the null object to every agent in* N^2 *and stop. Otherwise, continue to Step 2.*

Step $\boldsymbol{t \geqslant 2}$***.*** *Step* $t-1$ *ended with the definitions of* $N^t := N^{t-1} \backslash \bar{N}^{t-1} \neq \varnothing$ *and* $O^t := O^{t-1} \backslash \bar{O}^{t-1} \neq \varnothing$ *after the removal of top trading cycles in the (reduced) STTC graph* $\mathcal{G}^{t-1}$.

Use $N' = N^t$ *and* $O' = O^t$ *to construct the corresponding reduced instance (Definition 5.2) and STTC graph* $\mathcal{G}^t = \mathcal{G}(N', O', \omega', R', \mu', \rhd')$.

As in Step 1, identify all top trading cycles in $\mathcal{G}^t$ *and execute each cycle.*

- *Let* $\bar{N}^t$ *and* $\bar{O}^t$ *denote the sets of agents and objects, respectively, belonging to cycles in* $\mathcal{G}^t$*. Assign each cycle agent his corresponding allotment* $\mathrm{STTC}_i(R, \mu, \rhd)$*, remove all agents in* $\bar{N}^t$ *and all objects in* $\bar{O}^t$*, except for the null object, and set* $N^{t+1} := N^t \backslash \bar{N}^t$ *and* $O^{t+1} := O^t \backslash \bar{O}^t$.
- *If* $N^{t+1} = \varnothing$*, stop. If* $N^{t+1} \neq \varnothing$ *and* $O^{t+1} = \varnothing$*, assign the null object to every agent in* N^{t+1} *and stop. Otherwise, continue to Step* $t+1$.

Termination. *Since at least one agent is removed at each step, the procedure terminates after at most* n *steps, when each agent has received an allotment.*

Output: *Denote the resulting allocation by* $\mathrm{STTC}(R, \mu, \rhd)$.

An instance $(N, O, \omega, R, \mu, \rhd)$ *consists of a problem* (N, O, ω, R)*, an initial allocation* μ*, and a priority order* $\rhd$*. We refer to the STTC algorithm with priority order* $\rhd$ *and initial allocation* μ *as the* $(\boldsymbol{\rhd, \mu})$***-STTC algorithm****.*

Next, we introduce μ-cycles and non-μ-cycles. Along a μ-cycle, each agent points to the object he occupies at μ and receives it back, so no object changes hands. Thus, only non-μ-cycles can move objects between agents. The presence or absence of such cycles underlies the exclusion-core-stability test in Theorem 5.3.

[8] Note that each agent points to exactly one object, his most-preferred object, which is why we refer to the cycles that form as **top trading cycles**. At least one top trading cycle exists at every step: any agent pointing to $\varnothing$ forms the self-cycle $i \to \varnothing \to i$; if no agent points to $\varnothing$, then $\varnothing$ is isolated and every remaining vertex other than the isolated null object has out-degree one, so following the outgoing edges from any agent within the finite vertex set eventually revisits a vertex and thereby traces a cycle.

Definition 5.4 (μ- and non-μ-cycles) *Let $\mathcal{C}$ be a cycle in $\mathcal{G}^t$ at Step t of the $(\rhd, \mu)$-STTC algorithm, and let $S(\mathcal{C}) \subseteq N^t$ denote its set of agents. The cycle $\mathcal{C}$ is a **μ-cycle** if every $i \in S(\mathcal{C})$ points to $\mu(i)$; otherwise, it is a **non-μ-cycle**.*

*If a non-μ-cycle forms, let $t^\star$ be the first step at which one forms. Any non-μ-cycle in $\mathcal{G}^{t^\star}$ is called an **earliest non-μ-cycle**.*

Example D.1 in Appendix D illustrates the above notions.

5.3.2 STTC allocations and the exclusion core

The following proposition establishes that allocations produced by the $(\rhd, \mu)$-STTC algorithm are exclusion core stable.

Proposition 5.5 (STTC allocations are exclusion core stable) *For each problem R, each priority order $\rhd$, and each allocation μ, $\mathrm{STTC}(R, \mu, \rhd)$ is exclusion core stable.*

Given a priority order $\rhd$ and an allocation μ, the $(\rhd, \mu)$-STTC algorithm assigns each agent his most-preferred object among those available when he is assigned. Suppose, contrary to the proposition, that its outcome is exclusion-blocked. An induction over the steps shows that no agent assigned at a given step can belong to the blocking coalition and that no object removed at that step can enter the blocking coalition's extended eviction set. Consequently, no coalition member can obtain through exclusion-blocking an object he prefers to his STTC allotment, a contradiction. Appendix B contains the proof of Proposition 5.5.

When the initial allocation is the **null allocation $\mu^\varnothing := (\varnothing, \varnothing, \ldots, \varnothing)$**, the $(\rhd, \mu^\varnothing)$-STTC algorithm specializes to several well-known algorithms:

- For **Shapley–Scarf housing market problems**, each $(\rhd, \mu^\varnothing)$-STTC algorithm coincides with the classical **top trading cycles (TTC) algorithm** (L. Shapley & Scarf, 1974).
- For **house allocation problems**, the $(\rhd, \mu^\varnothing)$-STTC algorithm coincides with the **$\rhd$-serial dictatorship** algorithm (Svensson, 1999).
- For **house allocation problems with existing tenants**, the $(\rhd, \mu^\varnothing)$-STTC algorithm coincides with the **$\rhd$-you-request-my-house–I-get-your-turn ($\rhd$-YRMH-IGYT) algorithm** (Abdulkadiroğlu & Sönmez, 1999).

Next, we show that the exclusion core can be recovered from the outcomes of the $(\rhd, \mu^\varnothing)$-STTC algorithms. This result is closely related to Theorem 4 of Balbuzanov and Kotowski (2019); their Lemma 3(b) adapts that theorem to our model.[9]

[9] Balbuzanov and Kotowski (2019) allow object-specific priority orders, whereas our characterization uses a common priority order for all objects. This restriction strengthens the characterization but leaves less flexibility in the proof.

Proposition 5.6 (An exclusion core characterization) *For each problem R, the exclusion core is exactly the set of STTC outcomes obtained from the null allocation as the priority order varies over Δ:* $\mathrm{EC}(R) = \{\mathrm{STTC}(R, \mu^{\varnothing}, \rhd) \mid \rhd \in \Delta\}$.

By Proposition 5.5, every STTC outcome obtained from the null allocation belongs to the exclusion core. It remains to prove the reverse inclusion. Starting from an exclusion-core-stable allocation μ and an arbitrary priority order $\rhd^*$, we use the successive steps of the $(\rhd^*, \mu)$-STTC algorithm to construct a priority order $\rhd$ for which $\mathrm{STTC}(R, \mu^{\varnothing}, \rhd) = \mu$. Appendix B contains the proof of Proposition 5.6.[10]

Propositions 5.5 and 5.6 establish basic properties of the $(\rhd, \mu)$-STTC algorithm that are used to derive our main results in Section 5.4.

The next example illustrates Proposition 5.6. A more complex illustration appears in Example D.1 in Appendix D.

Example 5.7 (Example 5.9 continued) *Consider the problem from Example 5.9: $N = \{1, 2, 3, 4\}$, $O = \{a, b, c, d\}$; agents 1, 2, and 3 own objects a, b, and c, respectively, and d is publicly owned. Preferences are as follows.*

$$\begin{array}{lccccc} R_1: & b & a & \varnothing & \ldots & \\ R_2: & d & c & b & \varnothing & a \\ R_3: & b & c & \varnothing & \ldots & \\ R_4: & a & d & c & \varnothing & b. \end{array}$$

As $\rhd$ varies over Δ, the outcomes of the $(\rhd, \mu^{\varnothing})$-STTC algorithm are exactly the two allocations $\mu^1 = (a, d, b, c)$ and $\mu^2 = (b, d, c, a)$. $\mathrm{STTC}(R, \mu^{\varnothing}, \rhd) = \mu^1$ *if and only if agent 3 has the highest priority under $\rhd$, or agent 2 has the highest priority and agent 3 the second highest; otherwise* $\mathrm{STTC}(R, \mu^{\varnothing}, \rhd) = \mu^2$. *Hence, by Proposition 5.6, the exclusion core is $\{\mu^1, \mu^2\}$.* $\diamond$

Before turning to our main results, we relate our STTC algorithms to the GTTC algorithms of Balbuzanov and Kotowski (2019) and record two generalizations of Proposition 5.5.

The **generalized top trading cycles (GTTC) algorithm** of Balbuzanov and Kotowski (2019, Algorithm 2) includes our $(\rhd, \mu^{\varnothing})$-STTC algorithm as a special case. Their algorithm uses an object-specific priority structure, so different objects may have different priority orders. When $\mu \neq \mu^{\varnothing}$, our $(\rhd, \mu)$-STTC algorithms differ in form from their $\rhd$-GTTC algorithm. In the proof of their Proposition 4, however, Balbuzanov and Kotowski (2019) use a modification of their GTTC algorithm that corresponds to our

[10] Zhang (2026, Proposition 5) also characterizes the exclusion core on this domain, using a generalization of the YRMH-IGYT algorithm. Our maximal domain result (Corollary 5.2) and the results of Section 5.4, which rely on the STTC algorithm being initialized at an arbitrary allocation, have no counterpart there.

STTC algorithm. The correspondence between the two algorithms yields the following common generalization.

Remark 5.8 (Flexible TTC graphs, algorithms, and allocations) *In an STTC graph, the priority order $\triangleright$ is used only to determine the outgoing edge of each unassigned public object. These edges may be chosen arbitrarily rather than through a common priority order, for instance depending on the cycles executed at earlier steps.*

*We refer to the resulting constructions as **flexible top trading cycles (flexible TTC) graphs, algorithms**, and **allocations**. For mixed-ownership problems, the class of flexible TTC algorithms includes both our STTC algorithms and the corresponding GTTC algorithms of Balbuzanov and Kotowski (2019). The proof of Proposition 5.5 uses no property of the edge choices for unassigned public objects beyond the fact that each such object points to a remaining agent. Hence, every allocation generated by a flexible TTC algorithm is exclusion core stable.*

This observation concerns only Proposition 5.5. Proposition 5.6 and Theorem 5.3, by contrast, draw their strength from the single priority order of the STTC family: the former asserts the existence of a priority order under which a given allocation is generated, the latter characterizes exclusion core stability through the fixed points of the $(\triangleright, \mu)$-STTC algorithm. $\diamond$

Remark 5.9 (Flexible TTC allocations for complex endowments) *Recall that in object allocation problems with complex endowments, an object may be owned by any coalition of agents, not only by a singleton or by the grand coalition. We can extend the flexible TTC construction to such problems by requiring an object to point to one of its remaining owners whenever at least one owner remains; once no owner remains, its outgoing edge can be chosen arbitrarily among the remaining agents. A flexible TTC algorithm satisfying this requirement is **owner-respecting**: an object points to an owner whenever one remains, so owners' claims are not bypassed while they are present.*

*The argument of Proposition 5.5 extends to this class. Indeed, if a set $T \subsetneq N$ consists only of agents who remain at some step and an object belongs to the direct eviction set of T, then all owners of that object belong to T and hence remain; in particular, the object is not public, so owner-respectingness applies to it. Since all of its owners remain, they were present at every earlier step. If the object had been removed in an earlier cycle, owner-respectingness would have made it point to one of its owners, who would then have been removed with it, a contradiction. Thus, a coalition of agents who all remain at a given step has no exclusion rights over objects that have already been assigned, exactly as in the proof of Proposition 5.5, and the rest of that proof is unchanged. Hence, every allocation generated by an **owner-respecting flexible TTC algorithm** is exclusion core stable. For the special case in which each object's edge is determined by a fixed priority order over agents, this also follows from Corollary 1(a) and Lemma 3(a) of Balbuzanov and Kotowski (2019). Owner-respecting flexible TTC algorithms are more general: an object may point to different owners at different steps, as long as each is a remaining owner, and such edge choices need not be generated by any fixed order.* $\diamond$

5.4 Main Results

5.4.1 A maximal domain result

Proposition 5.6 raises the question of whether the STTC characterization of the exclusion core extends beyond mixed ownership. The next result shows that it does not when the characterization is required to hold for every preference profile. Every complex endowment structure containing a properly co-owned object admits a preference profile for which some exclusion-core-stable allocation cannot be obtained by the extension of the STTC algorithm defined below. In fact, that allocation cannot be obtained by any owner-respecting flexible TTC algorithm described in Remark 5.9.

Throughout this subsection, we allow proper co-ownership and write (ω, R) for a problem (N, O, ω, R). Accordingly, ω_i now denotes the set of objects that agent i owns, alone or as a proper co-owner. As before, public objects do not belong to any individual set ω_i. An object $a \in O$ is **properly co-owned** if at least two but not all agents own it; that is, there exist distinct agents $i, j \in N$ with $a \in \omega_i \cap \omega_j$ and an agent $k \in N$ with $a \notin \omega_k$. Under mixed ownership, by contrast, every object is owned either by exactly one agent or by all agents. For each object $o \in O$, let $S^o \subseteq N$ denote its **set of owners**, defined by

$$S^o := \begin{cases} \{i \in N \mid o \in \omega_i\}, & \text{if } \{i \in N \mid o \in \omega_i\} \neq \varnothing, \\ N, & \text{otherwise.} \end{cases}$$

The direct eviction set of a coalition S is $\omega(S) := \{o \in O \mid S^o \subseteq S\}$. We call the members of S^o the owners of o, and use *co-owner* only when o is properly co-owned and the presence of several owners matters. Thus, a coalition can directly evict others from an object exactly when it contains all of that object's owners.[11] The extended eviction set is obtained from this direct eviction set by the recursion introduced in Section 5.2.2.

In an extended $(\triangleright, \mu)$-STTC algorithm, a properly co-owned object points as follows. If its allottee at μ is a remaining co-owner, it points to that agent. Otherwise, if any co-owner remains, it points to the highest-priority remaining co-owner; if none remains, it points to the highest-priority remaining agent. The outgoing-edge rules for individually and publicly owned objects remain unchanged.

For a problem (ω, R), let $\mathrm{STTC}(\omega, R, \mu, \triangleright)$ denote the outcome of the extended $(\triangleright, \mu)$-STTC algorithm. The set of **STTC allocations generated from the null allocation** is

$$\mathrm{STTC}(\omega, R) := \{\mathrm{STTC}(\omega, R, \mu^{\varnothing}, \triangleright) \mid \triangleright \in \Delta\}.$$

We denote the exclusion core of the problem (ω, R) by $\mathrm{EC}(\omega, R)$.

[11] Whenever no object is properly co-owned, this coincides with the direct eviction set of Section 5.2.1: for $S \subsetneq N$, an individually owned object requires its single owner and a public object requires N, so $\omega(S) = \bigcup_{i \in S} \omega_i$; and $\omega(N) = O$.

Theorem 5.1 (Strict inclusion of the STTC under proper co-ownership) *Let ω be an endowment structure with a properly co-owned object. Then there exists a preference profile R such that* $\mathrm{STTC}(\omega, R) \subsetneq \mathrm{EC}(\omega, R)$.

The proof fixes a properly co-owned object a and chooses preferences under which every agent ranks a above $\varnothing$ and every other object below $\varnothing$. Let μ^a denote the allocation that assigns a to a non-co-owner and $\varnothing$ to every other agent. By Remark 5.9, for every priority order $\rhd$, $\mathrm{STTC}(\omega, R, \mu^{\varnothing}, \rhd) \in \mathrm{EC}(\omega, R)$. Hence, $\mathrm{STTC}(\omega, R) \subseteq \mathrm{EC}(\omega, R)$. The allocation μ^a is exclusion core stable but is not an STTC outcome. Therefore, $\mathrm{STTC}(\omega, R) \subsetneq \mathrm{EC}(\omega, R)$. Appendix C contains the proof of Theorem 5.1.

In fact, μ^a also lies outside the range of every owner-respecting flexible TTC algorithm described in Remark 5.9. The underlying tension is that these trading-cycles algorithms respect the co-owners' claims, whereas exclusion core stability may sustain an outsider's consumption of a properly co-owned object because a blocking coalition must make all its members strictly better off.

Combining Proposition 5.6 and Theorem 5.1 yields a maximal domain result.

Corollary 5.2 (A maximal domain result) *For any endowment structure ω, the following statements are equivalent:*

(i) *for every preference profile R,* $\mathrm{STTC}(\omega, R) = \mathrm{EC}(\omega, R)$*;*
(ii) *ω is a mixed-ownership structure, i.e., each object is either individually or publicly owned.*

Appendix C contains the proof of Corollary 5.2. Corollary 5.2 shows that, among domains defined by restrictions on endowment structures, the mixed-ownership domain is the largest on which the STTC algorithms characterize the exclusion core.

5.4.2 An exclusion-core-stability test

The STTC algorithm also provides a simple test for exclusion core stability. An allocation is exclusion core stable exactly when it is a fixed point of the $(\rhd, \mu)$-STTC algorithm—for one, and hence every, priority order $\rhd$. Otherwise, an earliest non-μ-cycle identifies an exclusion-blocking coalition.

Theorem 5.3 (An exclusion-core-stability test) *Let R be a problem and μ an allocation. Then the following statements are equivalent:*

(i) *μ is exclusion core stable;*
(ii) *for some priority order $\rhd$, $\mu = \mathrm{STTC}(R, \mu, \rhd)$;*
(iii) *for every priority order $\rhd$, $\mu = \mathrm{STTC}(R, \mu, \rhd)$.*

That (ii) implies (i) is Proposition 5.5, and that (iii) implies (ii) is immediate. For the remaining implication, suppose that μ is exclusion core stable and that $\mathrm{STTC}(R, \mu, \rhd) \neq \mu$ for some priority order $\rhd$. Then an earliest non-μ-cycle exists, and the agents in that cycle who strictly improve upon their μ-allotments form an exclusion-blocking coalition, contradicting the exclusion-core stability of μ. Appendix C contains the proof of Theorem 5.3. Example D.1 in Appendix D illustrates the earliest-non-μ-cycle argument used in the proof of Theorem 5.3.

The construction behind Theorem 5.1 also shows that this fixed-point characterization does not extend to object allocation problems with proper co-ownership. Recall the allocation μ^a defined after Theorem 5.1, which is exclusion core stable. However, when an owner-respecting flexible TTC algorithm is initialized at μ^a, the properly co-owned object is reassigned from a non-co-owner to a co-owner at the first step. Hence, the implication from exclusion core stability to the fixed-point condition fails.

5.4.3 Von Neumann–Morgenstern stability

We first introduce the standard definition of a von Neumann–Morgenstern stable set (1944) with respect to an arbitrary domination relation between allocations.

Definition 5.4 (Von Neumann–Morgenstern stable sets) *A set of allocations $\mathcal{V}$ is **von Neumann–Morgenstern stable** with respect to a domination relation if it is*

- ***internally stable**: for all distinct $\mu, \nu \in \mathcal{V}$, ν does not dominate μ; and*
- ***externally stable**: for each $\mu \notin \mathcal{V}$, there exists $\nu \in \mathcal{V}$ that dominates μ.*

The domination relation we use is the one induced by exclusion-blocking.

Definition 5.5 (Exclusion-domination) *Given a problem R and allocations μ and ν, we say that ν **exclusion-dominates** μ if some coalition S exclusion-blocks μ via ν.*

We next show that the exclusion core is always a von Neumann–Morgenstern stable set with respect to exclusion-domination. In addition, we show that, for every allocation $\mu \notin \mathrm{EC}(R)$ and every priority order $\rhd$, the outcome of the $(\rhd, \mu)$-STTC algorithm belongs to the exclusion core and exclusion-dominates μ.

Theorem 5.6 (Von Neumann–Morgenstern stability of the exclusion core) *Let R be a problem. Then, $\mathrm{EC}(R)$ is a von Neumann–Morgenstern stable set with respect to exclusion-domination, i.e., it is*

- ***internally stable**: for all distinct $\mu, \nu \in \mathrm{EC}(R)$, ν does not exclusion-dominate μ; and*
- ***externally stable**: for each $\mu \notin \mathrm{EC}(R)$ and each priority order $\rhd$, $\mathrm{STTC}(R, \mu, \rhd)$ belongs to $\mathrm{EC}(R)$ and exclusion-dominates μ.*

We verify the two stability conditions separately.

Internal stability. Internal stability follows directly from the definition of the exclusion core: no allocation exclusion-dominates an exclusion-core-stable allocation.

External stability. External stability is the substantive condition, and Theorem 5.6 asserts more than the existence of some dominating exclusion-core-stable allocation: for each priority order $\triangleright$, the allocation $\mathrm{STTC}(R, \mu, \triangleright)$ both belongs to the exclusion core and exclusion-dominates μ. The coalition that carries out this exclusion-blocking is assembled from the cycles of the $(\triangleright, \mu)$-STTC algorithm, and the same cycles generate the blocking path of Theorem 5.11. We therefore prove the two theorems together: the joint argument shows how the same STTC cycles generate both local blocking transitions and a global exclusion-dominating allocation. Section 5.4.4 outlines the common construction, Appendix C contains the joint proof, and Example D.1 in Appendix D illustrates both conclusions on a concrete problem.

Balbuzanov and Kotowski (2019, Proposition 4) establish a related stability result for what they call relational economies with acyclic priority structures. Their proposition concerns the strong and weak exclusion cores of such an economy, which are derived from conditional endowments that vary with the prevailing allocation; by their Lemma 3(b), these solutions coincide with the exclusion core of an object allocation problem exactly when every object is individually or publicly owned. How far Theorem 5.6 itself extends is taken up in Appendix E. An endowment structure is acyclic in the sense of Balbuzanov and Kotowski (2019) exactly when either no object is properly co-owned or all objects share a single set of owners (Proposition E.2 in Appendix E). Theorem E.5 (Appendix E) establishes von Neumann–Morgenstern stability on the second class, so the conclusion of Theorem 5.6 holds throughout the acyclic class (Corollary E.7 in Appendix E). Beyond the acyclic class, external stability may fail: Appendix E Example E.9 exhibits a complex-endowment problem whose exclusion core is not externally stable,[12] whereas Appendix E Example E.10 exhibits a non-acyclic endowment structure whose exclusion core is externally stable at every preference profile.

5.4.4 Paths to exclusion-core stability

The core concepts introduced in Section 5.2.2 are static: they are defined by the absence of blocking. Blocking, however, describes a coalitional deviation and is inherently dynamic. This raises a natural question: does an explicit process of blocking lead to core allocations? Serrano and Volij (2008) answer this question affirmatively for classical core concepts.

Paths to stability have also been established for marriage markets by Roth and Vande Vate (1990), for roommate markets admitting stable matchings by Diamantoudi et al.

[12] We are grateful to Jacob Coreno for this example.

(2004), and for matching markets with couples under weak responsiveness by Klaus and Klijn (2007). We now tackle this question for exclusion-blocking and the exclusion core.

When a coalition S exclusion-blocks μ via σ, every member of S is strictly better off at σ. Implementing σ, however, may require agents outside S who also benefit from it. Thus, S need not have sufficient exclusion rights to implement the blocking on its own. The following simple example illustrates this point.

Suppose that $N = \{1, 2\}$, agent 1 owns object a, agent 2 owns object b, and preferences satisfy $b\ P_1\ a\ P_1\ \varnothing$ and $a\ P_2\ b\ P_2\ \varnothing$. Consider $\mu = (a, b)$ and $\sigma = (b, a)$. Allocation σ Pareto dominates μ, and coalition $\{1\}$ exclusion-blocks μ via σ: agent 1 is strictly better off, while agent 2 is also better off and therefore need not be excluded. However, if only agent 1 participates and receives object b, the resulting allocation is $(b, \varnothing)$. Since $b \notin \Omega(\{1\} \mid \mu) = \{a\}$, coalition $\{1\}$ cannot exclusion-block μ via this allocation. Thus, although $\{1\}$ exclusion-blocks μ via σ, it cannot implement this blocking on its own.

To model exclusion-blocking dynamically, we distinguish the blocking allocation σ from the allocation obtained when coalition members receive their σ-allotments and outsiders' allotments are disrupted as little as possible. This distinction leads to the notions of post-blocking allocations, self-contained exclusion-blockings, and the paths they generate.

Definition 5.7 (Post-blocking and self-contained exclusion-blockings) *Given a problem R, allocations μ and σ, and a coalition S that exclusion-blocks μ via σ, the **post-blocking allocation induced by** (S, σ) **at** μ, denoted by $\widehat{\sigma}^{S,\mu}$, is defined by*

$$\widehat{\sigma}^{S,\mu}(i) = \begin{cases} \sigma(i) & \text{if } i \in S, \\ \varnothing & \text{if } i \notin S \text{ and } \mu(i) \in \sigma(S)\backslash\{\varnothing\}, \\ \mu(i) & \text{otherwise.}^{13} \end{cases} \tag{5.1}$$

*We also say that $\widehat{\sigma}^{S,\mu}$ is obtained by satisfying the exclusion-blocking of μ by S via σ. This exclusion-blocking is **self-contained** if S also exclusion-blocks μ via $\widehat{\sigma}^{S,\mu}$.*

We use self-contained exclusion-blockings to define the corresponding dynamic paths.

Definition 5.8 (Self-contained exclusion-blocking paths) *For $K \geqslant 0$, a sequence of allocations $\mu^0 \mapsto \mu^1 \mapsto \ldots \mapsto \mu^K$ is **a self-contained exclusion-blocking path of length K from μ^0 to μ^K** if, for every $\ell \in \{0, 1, \ldots, K-1\}$, $\mu^{\ell+1}$ is the post-blocking allocation induced by a self-contained exclusion-blocking of μ^ℓ.*

[13] Since the null object has enough capacity for all agents, an outsider whose μ-allotment is $\varnothing$ is never displaced; hence, the eviction clause refers to $\sigma(S)\backslash\{\varnothing\}$. The resulting map is an allocation: no real object is assigned to more than one coalition member; an outsider retains a real μ-allotment only if no coalition member receives it; and all displaced outsiders receive the null object, which has enough capacity for all agents.

The next lemma shows that restricting attention to self-contained exclusion-blockings does not eliminate any blocking opportunity.

Lemma 5.9 (Self-contained completion of an exclusion-blocking coalition) *Suppose S exclusion-blocks μ via σ, and let $S^+(\mu,\sigma) := \{i \in N \mid \sigma(i)\ P_i\ \mu(i)\}$ be the set of agents who are strictly better off at σ than at μ. Then $S \subseteq S^+(\mu,\sigma)$, $S^+(\mu,\sigma)$ exclusion-blocks μ via σ, and this exclusion-blocking is self-contained.*

Appendix C contains the proof of Lemma 5.9.

In the preceding illustration, $S^+(\mu,\sigma) = \{1,2\}$, and its induced post-blocking allocation is σ. The completion therefore adds the improving outsider whose participation is needed to implement the blocking. The set $S^+(\mu,\sigma)$ contains every agent who improves, so it is the largest coalition that can carry out this blocking; in general, it need not be inclusion-minimal among the self-contained blocking coalitions.

Because self-contained exclusion-blockings capture every blocking opportunity and each of them induces a well-defined successor allocation, they can be applied repeatedly. Whether doing so leads to the exclusion core is the central question of the dynamic approach: does every allocation admit a self-contained exclusion-blocking path (Definition 5.8) to the exclusion core? We show that the answer is affirmative.

Before presenting our convergence result, we show that arbitrary sequences of self-contained exclusion-blockings need not converge. Small blocking coalitions may be particularly likely to form. An agent can identify a profitable unilateral deviation without coordinating with others, and two agents may discover a joint deviation through bilateral communication. By contrast, forming and coordinating a larger coalition may be more difficult. Motivated by this possibility, in the following example we repeatedly execute the smallest self-contained exclusion-blocking coalitions, which have size one or two.

Example 5.10 (The self-contained exclusion-blocking dynamics can cycle) *Let $N = \{1,2,3\}$, $O = \{a,b,c\}$, and endowments $\omega_1 = \{a\}$, $\omega_2 = \{b\}$, and $\omega_3 = \{c\}$. Preferences are given by*

$$\begin{array}{lllll} R_1: & b & c & a & \varnothing \\ R_2: & c & a & b & \varnothing \\ R_3: & a & b & c & \varnothing. \end{array}$$

Consider the following sequence of allocations, obtained by successively satisfying self-contained exclusion-blockings of minimum size. Starting from $\mu^1 = (b,a,c)$, the successive blockings are summarized in the following table.

Step	*Blocking coalition and gains*	*Resulting allocation*	*Displaced outsider*
1	$\{2,3\}$: $c\ P_2\ a$, $b\ P_3\ c$	$\mu^2 = (\varnothing, c, b)$	*Agent* 1, *from* $b \in \omega_2$
2	$\{1\}$: $a\ P_1\ \varnothing$	$\mu^3 = (a, c, b)$	*None*
3	$\{1,3\}$: $c\ P_1\ a$, $a\ P_3\ b$	$\mu^4 = (c, \varnothing, a)$	*Agent* 2, *from* $c \in \omega_3$
4	$\{2\}$: $b\ P_2\ \varnothing$	$\mu^5 = (c, b, a)$	*None*
5	$\{1,2\}$: $b\ P_1\ c$, $a\ P_2\ b$	$\mu^6 = (b, a, \varnothing)$	*Agent* 3, *from* $a \in \omega_1$
6	$\{3\}$: $c\ P_3\ \varnothing$	$\mu^1 = (b, a, c)$	*None*

At each step, every member of the listed coalition is strictly better off. At Steps 1, 3, and 5, the only displaced outsider occupies an object owned by a member of the blocking coalition; at Steps 2, 4, and 6, no outsider is displaced. Hence, every transition is generated by a self-contained exclusion-blocking, and Step 6 returns to μ^1, so the sequence cycles.

At every step, the listed coalition has minimum size among the available self-contained exclusion-blocking coalitions. At Steps 2, 4, and 6, they are singletons.

At Step 1, agent 1 already receives his most-preferred object, while $c \notin \Omega(\{2\} \mid \mu^1) = \{a, b\}$ and $\{a, b\} \cap \Omega(\{3\} \mid \mu^1) = \varnothing$, so neither agent 2 nor agent 3 can block alone. Similarly, at Step 3, $\{b, c\} \cap \Omega(\{1\} \mid \mu^3) = \varnothing$ and $a \notin \Omega(\{3\} \mid \mu^3) = \{b, c\}$, and at Step 5, $b \notin \Omega(\{1\} \mid \mu^5) = \{a, c\}$ and $\{a, c\} \cap \Omega(\{2\} \mid \mu^5) = \varnothing$. Thus, no singleton exclusion-blocks at Steps 1, 3, or 5.

Allocation $\mu = (b, c, a)$ assigns every agent his most-preferred object. It therefore Pareto dominates every other allocation and is the unique Pareto efficient allocation. By Lemma 5.7, for each $k \in \{1, \ldots, 6\}$, the agents who are strictly better off at μ than at μ^k exclusion-block μ^k via μ. $\diamond$

Although the exclusion-blocking dynamics can cycle, the next theorem shows that every allocation admits some self-contained exclusion-blocking path to the exclusion core.

Theorem 5.11 (Paths to exclusion-core stability) *Let R be a problem and μ an allocation. Then there exist $K \geqslant 0$ and allocations $\mu^0, \mu^1, \ldots, \mu^K$ such that $\mu^0 = \mu$, μ^K is exclusion core stable, and $\mu^0 \mapsto \mu^1 \mapsto \cdots \mapsto \mu^K$ is a self-contained exclusion-blocking path.*

Appendix C contains the joint proof of Theorems 5.6 and 5.11. We outline the common construction here.

Let μ be an allocation that is not exclusion core stable, fix a priority order $\rhd$, and set $\mu^{\rhd} := \mathrm{STTC}(R, \mu, \rhd)$. By Proposition 5.5, $\mu^{\rhd}$ is exclusion core stable, and by Theorem 5.3, $\mu^{\rhd} \neq \mu$.

Follow the steps of the $(\rhd, \mu)$-STTC algorithm at which non-μ-cycles form. At each such step, assign the cycle agents their final STTC allotments and assign the null

object to every outsider whose current allotment is taken. Let the cycle agents who strictly improve relative to the current allocation form the local coalition. Whenever this coalition is nonempty, a backward ownership-chain argument shows that it exclusion-blocks the current allocation and that the resulting allocation is precisely its post-blocking allocation. If the coalition is empty, the constructed allocation is unchanged. Omitting these repetitions yields a self-contained exclusion-blocking path from μ to $\mu^{\triangleright}$.

The same induction establishes external stability. The agents encountered along the construction who prefer their final STTC allotments to their original allotments form a cumulative coalition that exclusion-blocks μ via $\mu^{\triangleright}$.

Example D.1 in Appendix D illustrates the construction.

How far Theorem 5.11 extends beyond mixed ownership is taken up in Appendix E. There, Proposition E.2 shows that an endowment structure is acyclic exactly when either no object is properly co-owned or all objects share a single set of owners. The first case is covered by Theorem 5.11 itself and the second by Theorem E.6 (Appendix E), so the conclusion of Theorem 5.11 holds throughout the acyclic class (Corollary E.8 in Appendix E).

Theorem 5.11 exhibits one path to the exclusion core, constructed from the cycles of a single STTC algorithm. A decentralized process need not follow that path: at an allocation outside the exclusion core, several self-contained exclusion-blockings may be available, and which of them is carried out is not determined by the model. We therefore ask what happens when the blocking that occurs at each step is left to chance. The following process makes this precise.

Fix a problem R and an allocation μ, and set $\mu^1 := \mu$. Given μ^ℓ, if μ^ℓ is exclusion core stable, set $\mu^{\ell+1} := \mu^\ell$. Otherwise, choose at random a post-blocking allocation induced by a self-contained exclusion-blocking of μ^ℓ and denote it by $\mu^{\ell+1}$.

The set of available post-blocking allocations is nonempty whenever μ^ℓ is not exclusion core stable. Indeed, some coalition then exclusion-blocks μ^ℓ, and Lemma 5.9 yields a self-contained completion and hence an available post-blocking allocation. Thus, the process is well defined.

Theorem 5.11 yields the following corollary.

Corollary 5.12 (Random paths to exclusion-core stability) *Let R be a problem and μ an allocation. Suppose that, in the random process defined above, every available post-blocking allocation is selected with positive probability and the selection probabilities depend only on the current allocation. Then, with probability one, the sequence $(\mu^\ell)_{\ell \geqslant 1}$ reaches an exclusion-core-stable allocation after finitely many steps and remains there.*

Appendix C contains the proof of Corollary 5.12.

The next example applies the path result to Example 5.6 and clarifies the distinction between external stability and dynamic reachability. External stability is a static domination property, whereas a self-contained exclusion-blocking path describes how an exclusion-core-stable allocation can be reached through successive deviations.

Example 5.13 (Example 5.6 continued) *Consider again the problem from Example 5.6, with* $N = \{1, 2, 3\}$ *and* $O = \{a, b, c\}$. *Agents* 1, 2, *and* 3 *own objects* a, b, *and* c, *respectively. Preferences are as follows.*

$$\begin{array}{llllll} R_1: & \underline{a} & \boldsymbol{c} & b & \varnothing \\ R_2: & \boldsymbol{a} & \underline{b} & c & \varnothing \\ R_3: & \boldsymbol{b} & a & \underline{c} & \varnothing \end{array}$$

Recall that coalition $S = \{1\}$ *exclusion-blocks allocation* $\boldsymbol{\mu} = (\boldsymbol{c}, \boldsymbol{a}, \boldsymbol{b})$ *via the exclusion-core-stable allocation* $\underline{\sigma} = (\underline{a}, \underline{b}, \underline{c})$. *At* μ, $\{1\}$ *is the only coalition that exclusion-blocks* μ. *We construct a self-contained exclusion-blocking path of length three from* μ *to* $\underline{\sigma}$ *and then show that no shorter path reaches the exclusion core. The three transitions are summarized in the following table.*

Step	*Blocking coalition and gains*	*Resulting allocation*	*Displaced outsider*
1	$\{1\}$: $a\ P_1\ c$	$\mu^1 = (a, \varnothing, b)$	*Agent* 2, *from* $a \in \omega_1$
2	$\{2\}$: $b\ P_2\ \varnothing$	$\mu^2 = (a, b, \varnothing)$	*Agent* 3, *from* $b \in \omega_2$
3	$\{3\}$: $c\ P_3\ \varnothing$	$\underline{\sigma} = (a, b, c)$	*None*

At each step, the listed coalition exclusion-blocks the preceding allocation via the resulting post-blocking allocation. Hence, $\mu \mapsto \mu^1 \mapsto \mu^2 \mapsto \underline{\sigma}$ *is a self-contained exclusion-blocking path to the exclusion core.*

Minimality. *No self-contained exclusion-blocking path from* μ *to the exclusion core has length less than three. At* μ, *agents* 2 *and* 3 *receive their most-preferred objects, so only agent* 1 *can be made strictly better off; hence* $S = \{1\}$, *and agent* 1 *can improve only by receiving object* a. *The induced post-blocking allocation is therefore* μ^1, *the unique successor of* μ. *At* μ^1, *agents* 1 *and* 3 *receive their most-preferred objects, so any exclusion-blocking coalition must be* $\{2\}$. *Since* $\Omega(\{2\} \mid \mu^1) = \{b, c\}$, *agent* 2 *can obtain only* b *or* c. *These choices yield, respectively,* μ^2 *and* (a, c, b). *Neither equals* $\underline{\sigma}$, *and* $\mathrm{EC}(R) = \{\underline{\sigma}\}$ *by Proposition 5.6. Hence, no self-contained exclusion-blocking path of length at most two reaches the exclusion core.*

Although $\{1\}$ *exclusion-blocks* μ *via* σ, *satisfying this blocking under the minimally disruptive rule yields* μ^1, *not* σ. *The blocking is self-contained because* $\{1\}$ *also exclusion-blocks* μ *via* μ^1, *but reaching* σ *requires the three-step path above. Thus, external stability does not itself provide a dynamic route to the exclusion core.* ◇

5.5 Concluding Remarks

We study the exclusion core in mixed-ownership object allocation problems and introduce the serial top trading cycles algorithm as a tool for analyzing it. For every initial allocation and every priority order, the STTC algorithm produces an exclusion-core-stable allocation. When the STTC algorithm is initialized at the null allocation, the set of its outcomes across priority orders coincides with the exclusion core. A single properly co-owned object is enough for this coincidence to fail. Among domains defined by restrictions on endowment structures, mixed ownership is therefore the largest domain on which the STTC algorithms generate exactly the exclusion core. The algorithm also provides a fixed-point test for exclusion core stability.

The sequence of cycles generated by the STTC algorithm provides a common constructive basis for two further stability results. Locally, the agents who improve at successive non-μ-cycle steps generate a self-contained exclusion-blocking path to the exclusion core. Globally, the agents who participate in non-μ-cycles and prefer their final STTC allotments to their initial allotments form a cumulative coalition that exclusion-blocks the initial allocation. The path result also has a dynamic consequence. When the blocking carried out at each step is left to chance—with selection probabilities that depend only on the current allocation and give every available post-blocking allocation positive probability—the process reaches the exclusion core in finitely many steps with probability one.

Proper co-ownership marks a sharp boundary for the characterization and the fixed-point test. In the presence of a single properly co-owned object, an exclusion-core-stable allocation may lie outside the range of every owner-respecting flexible TTC algorithm, and the fixed-point test may fail. The two stability results reach further. Appendix E characterizes the acyclic endowment structures, those in which either no object is properly co-owned or all objects share a single set of owners, and shows that both results continue to hold on this larger class. Acyclicity is not necessary for external stability either: beyond it the exclusion core may fail to be externally stable, but it may also be externally stable at every preference profile. Whether every allocation admits a self-contained exclusion-blocking path to the exclusion core in the presence of proper co-ownership remains an open question; in both complex-endowment examples of Appendix E such a path does exist.

Appendices

Appendix A: Proofs of the results in Section 5.2

Proof. **[Proof of Lemma 5.5]** Fix a problem R, an allocation μ, and coalitions S and S'. Recall from Section 5.2.2 that $S_0^\mu := S$ and that, for each $k \geqslant 0$, $S_{k+1}^\mu := S_k^\mu \cup \mu^{-1}\big(\omega(S_k^\mu)\big)$. This sequence is increasing, and since N is finite it is constant from some index on. Let ℓ be the smallest such index, so that $\Omega(S \mid \mu) = \omega(S_\ell^\mu)$. The four parts assert, respectively, that $\Omega(S \mid \mu)$ contains the direct eviction set, is monotone in the coalition, is closed under the endowments of agents whose allotments it contains, and propagates along chains of dependence.

We first record that the endowment function ω is **monotone**: for all coalitions $T \subseteq T'$, $\omega(T) \subseteq \omega(T')$. (This is distinct from the monotonicity of $\Omega(\cdot \mid \mu)$ asserted in part (b).) If $T' = N$, then $\omega(T') = O \supseteq \omega(T)$. If $T' \subsetneq N$, then $\omega(T) = \bigcup_{i \in T} \omega_i \subseteq \bigcup_{i \in T'} \omega_i = \omega(T')$.

(a) Since $S = S_0^\mu \subseteq S_\ell^\mu$, by monotonicity of ω, $\omega(S) \subseteq \omega(S_\ell^\mu) = \Omega(S \mid \mu)$.

(b) Suppose that $S \subseteq S'$, and let ℓ' be the smallest index from which $(S_k'^\mu)_{k \geqslant 0}$ is constant. We show by induction that, for every $k \geqslant 0$, $S_k^\mu \subseteq S_k'^\mu$. For $k = 0$, $S_0^\mu = S \subseteq S' = S_0'^\mu$. For the induction step, let $k \geqslant 0$ and suppose that $S_k^\mu \subseteq S_k'^\mu$. Then, by monotonicity of ω, $\omega(S_k^\mu) \subseteq \omega(S_k'^\mu)$. Hence, $\mu^{-1}(\omega(S_k^\mu)) \subseteq \mu^{-1}(\omega(S_k'^\mu))$. Thus, $S_{k+1}^\mu \subseteq S_{k+1}'^\mu$. Setting $\bar{k} := \max\{\ell, \ell'\}$, we obtain $\Omega(S \mid \mu) = \omega(S_{\bar{k}}^\mu) \subseteq \omega(S_{\bar{k}}'^\mu) = \Omega(S' \mid \mu)$.

(c) Let $i \in N$ and suppose that $\mu(i) \in \Omega(S \mid \mu) = \omega(S_\ell^\mu)$. Then $i \in \mu^{-1}(\omega(S_\ell^\mu)) \subseteq S_{\ell+1}^\mu = S_\ell^\mu$, and therefore $\omega_i \subseteq \omega(S_\ell^\mu) = \Omega(S \mid \mu)$.

(d) Let $q \geqslant 1$, let agents $i^1, \ldots, i^q$ and real objects $o^1, \ldots, o^q$ form a chain of dependence, and suppose that S reaches i^q; that is, $i^q \in S$ or $\omega_{i^q} \subseteq \Omega(S \mid \mu)$. We show by downward induction on r that, for each $r \in \{1, \ldots, q\}$, $o^r \in \Omega(S \mid \mu)$. Let $r = q$. If $i^q \in S$, then by (a), $o^q \in \omega_{i^q} \subseteq \omega(S) \subseteq \Omega(S \mid \mu)$. If $\omega_{i^q} \subseteq \Omega(S \mid \mu)$, then $o^q \in \omega_{i^q} \subseteq \Omega(S \mid \mu)$. Now let $r < q$ and suppose that $o^{r+1} \in \Omega(S \mid \mu)$. Since $\mu(i^r) = o^{r+1}$, by part (c), $\omega_{i^r} \subseteq \Omega(S \mid \mu)$. Since $o^r \in \omega_{i^r}$, it follows that $o^r \in \Omega(S \mid \mu)$. In particular, $o^1 \in \Omega(S \mid \mu)$.

Proof. **[Proof of Lemma 5.7]** Let R be a problem and let μ and σ be allocations such that σ Pareto dominates μ. Define $\widehat{S} := \{i \in N \mid \sigma(i) \; P_i \; \mu(i)\}$. Since σ Pareto dominates μ, $\widehat{S} \neq \varnothing$. Moreover, every agent in $\widehat{S}$ is strictly better off at σ. Since no agent outside $\widehat{S}$ is worse off at σ, the exclusion-enforcement condition is vacuously satisfied. Therefore, $\widehat{S}$ exclusion-blocks μ via σ.

Appendix B: Proofs of the results in Section 5.3

Proof. [Proof of Proposition 5.5] Fix a problem R, a priority order $\rhd$, and an initial allocation μ, and set $\mu^{\rhd} := \mathrm{STTC}(R, \mu, \rhd)$. Suppose, contrary to the proposition, that a coalition S exclusion-blocks $\mu^{\rhd}$ via an allocation σ. Then:

1. each agent in S is strictly better off at σ, i.e., for each $i \in S$, $\sigma(i) \; P_i \; \mu^{\rhd}(i)$; and
2. if, for some $j \notin S$, $\mu^{\rhd}(j) \; P_j \; \sigma(j)$, then $\mu^{\rhd}(j) \in \Omega(S \mid \mu^{\rhd})$.

We first show that a coalition consisting only of agents who are still present at Step t has no exclusion rights over any object that has already been assigned. A present agent's endowment cannot have been removed while he remains, so the recursion defining the extended eviction set never leaves the set of agents present at Step t. Formally, for every Step t at which the remaining sets N^t and O^t are defined, let $O^{<t}$ denote the set of real objects removed before Step t; that is, $O^{<t} := O \backslash O^t$. We show that, for every set of agents $T \subseteq N^t$,

$$\Omega(T \mid \mu^{\rhd}) \cap O^{<t} = \varnothing. \tag{5.2}$$

Since $O^{<1} = \varnothing$, the claim is immediate for $t = 1$. Consider $t \geqslant 2$. Write $T_0^{\mu^{\rhd}} := T$ and $T_{k+1}^{\mu^{\rhd}} := T_k^{\mu^{\rhd}} \cup (\mu^{\rhd})^{-1}\big(\omega(T_k^{\mu^{\rhd}})\big)$ for the recursion of Section 5.2.2 applied to T at $\mu^{\rhd}$. We show by induction on k that every agent in this recursion remains present at Step t: for every $k \geqslant 0$, $T_k^{\mu^{\rhd}} \subseteq N^t$. Since $T_0^{\mu^{\rhd}} = T \subseteq N^t$, this is true for $k = 0$.

Suppose that $T_k^{\mu^{\rhd}} \subseteq N^t$. Since N^t is a proper subset of N, the direct eviction set $\omega(T_k^{\mu^{\rhd}})$ consists only of objects individually owned by agents in $T_k^{\mu^{\rhd}}$. Every such object remains available at Step t. Indeed, while its owner remains, the object points to that owner, so it could not have been removed in an earlier cycle without its owner also leaving in that cycle.

Consequently, any agent who receives an object in $\omega(T_k^{\mu^{\rhd}})$ at the final allocation $\mu^{\rhd}$ also leaves no earlier than Step t. Hence, $(\mu^{\rhd})^{-1}\big(\omega(T_k^{\mu^{\rhd}})\big) \subseteq N^t$, and therefore $T_{k+1}^{\mu^{\rhd}} \subseteq N^t$. This completes the induction.

The sequence $(T_k^{\mu^{\rhd}})_{k \geqslant 0}$ is increasing, and since N is finite it is constant from some index on. Let ℓ_T be the smallest such index. All agents in the terminal set $T_{\ell_T}^{\mu^{\rhd}}$ are therefore still present at Step t. Every object in $\Omega(T \mid \mu^{\rhd}) = \omega(T_{\ell_T}^{\mu^{\rhd}})$ is individually owned by one of these agents and hence also remains available at Step t. This proves (5.2).

By going through the steps of the $(\rhd, \mu)$-STTC algorithm, we show that agents receiving their STTC allotments are not members of the exclusion-blocking coalition S.

$(\rhd, \mu)$-STTC algorithm:

Step 1. At Step 1, every assigned agent receives his most-preferred object in $O \cup \{\varnothing\}$. Hence, no such agent is strictly better off at σ than at $\mu^{\triangleright}$, and none belongs to S. Hence, $S \subseteq N^2$. Applying (5.2) with $T = S$ and $t = 2$ shows that no object removed at Step 1 belongs to $\Omega(S \mid \mu^{\triangleright})$. In particular, if agent j is assigned his $\mu^{\triangleright}$-allotment at Step 1, then $\mu^{\triangleright}(j) \notin \Omega(S \mid \mu^{\triangleright})$. We proceed inductively.

Step $t \geqslant 2$. Recall that no agent receiving his $\mu^{\triangleright}$-allotment at Steps $1, \ldots, t-1$ can belong to the exclusion-blocking coalition S. Hence, coalition S is a subset of agents remaining after Step $t-1$, i.e., $S \subseteq N^t$. Applying (5.2) with $T = S$ shows that no object removed before Step t belongs to $\Omega(S \mid \mu^{\triangleright})$. Thus, if j is assigned his $\mu^{\triangleright}$-allotment at one of Steps $1, \ldots, t-1$, then $\mu^{\triangleright}(j) \notin \Omega(S \mid \mu^{\triangleright})$.

At Step t, every assigned agent receives his most-preferred object in $O^t \cup \{\varnothing\}$. We show that no such agent is strictly better off at σ. Suppose, to the contrary, that there exists an agent $i \in N^t$ who is assigned object $\mu^{\triangleright}(i)$ at Step t and satisfies $\sigma(i) \; P_i \; \mu^{\triangleright}(i)$.

Set $i^0 := i$ and $\tau_0 := t$. We construct recursively agents $i^0, i^1, \ldots$ and steps $\tau_0 > \tau_1 > \cdots$ at which they are assigned. Suppose that i^k and τ_k have been constructed. If $\mu^{\triangleright}(i^k) \; R_{i^k} \; \sigma(i^k)$, the construction stops. Otherwise $\sigma(i^k) \; P_{i^k} \; \mu^{\triangleright}(i^k)$; since i^k receives his most-preferred object in $O^{\tau_k} \cup \{\varnothing\}$, the object $\sigma(i^k)$ was not available at Step τ_k, so it is a real object assigned at some earlier Step $\tau_{k+1} < \tau_k$, say to agent i^{k+1}, and $\mu^{\triangleright}(i^{k+1}) = \sigma(i^k)$. Since the steps are positive integers and strictly decreasing, the construction stops after finitely many rounds, say at index m, with $\mu^{\triangleright}(i^m) \; R_{i^m} \; \sigma(i^m)$. By assumption $\sigma(i^0) \; P_{i^0} \; \mu^{\triangleright}(i^0)$, so the construction does not stop at $k = 0$; hence $m \geqslant 1$.

By construction, $\mu^{\triangleright}(i^m) = \sigma(i^{m-1})$, which is a real object. Since σ is an allocation, this object cannot also be assigned to i^m at σ. Hence, $\sigma(i^m) \neq \mu^{\triangleright}(i^m)$. Since P_{i^m} is the strict part of R_{i^m}, it follows that $\mu^{\triangleright}(i^m) \; P_{i^m} \; \sigma(i^m)$.

Agent i^m received his STTC allotment at a step strictly before Step t. By the induction hypothesis, $i^m \notin S$ and $\mu^{\triangleright}(i^m) \notin \Omega(S \mid \mu^{\triangleright})$. However, since coalition S exclusion-blocks $\mu^{\triangleright}$ via σ and $\mu^{\triangleright}(i^m) \; P_{i^m} \; \sigma(i^m)$, we must have $\mu^{\triangleright}(i^m) \in \Omega(S \mid \mu^{\triangleright})$, a contradiction.

We conclude that no Step t agent can be strictly better off at allocation σ, and therefore no agent receiving his $\mu^{\triangleright}$-allotment at Step t can belong to the exclusion-blocking coalition S. Hence, $S \subseteq N^{t+1}$. Applying (5.2) with $T = S$ and the remaining sets N^{t+1} and O^{t+1} shows that no object removed during Steps $1, \ldots, t$ belongs to $\Omega(S \mid \mu^{\triangleright})$. In particular, if agent j is assigned his $\mu^{\triangleright}$-allotment during one of these steps, then $\mu^{\triangleright}(j) \notin \Omega(S \mid \mu^{\triangleright})$.

Agents remaining at termination. Suppose the algorithm stops after Step t. If $N^{t+1} = \varnothing$, then $S \subseteq N^{t+1} = \varnothing$, contradicting that S is a coalition.

Suppose instead that $N^{t+1} \neq \varnothing$. Then $O^{t+1} = \varnothing$, and each remaining agent $i \in N^{t+1}$ is assigned $\mu^{\triangleright}(i) = \varnothing$, his most-preferred object in $O^{t+1} \cup \{\varnothing\} = \{\varnothing\}$. If some remaining agent i belonged to S, then $\sigma(i)\ P_i\ \varnothing$, so $\sigma(i)$ would be a real object. Since $O^{t+1} = \varnothing$, this object was removed at one of Steps $1, \ldots, t$ and assigned to an agent at that step. The chain argument from Step t, applied with $i^0 := i$ and $\tau_0 := t + 1$, again yields a contradiction.

Every agent is either assigned at one of the steps of the $(\triangleright, \mu)$-STTC algorithm or remains when the algorithm terminates. The preceding arguments show that neither type of agent can belong to S. Hence, $S = \varnothing$, contradicting that S is a coalition.

The following auxiliary lemma isolates an argument used in the proof of Lemma B.2 and twice in the joint proof of Theorems 5.6 and 5.11. It converts an STTC cycle into a dependence chain of the type required by Lemma 5.5(d).

The idea is to walk around the cycle in the direction of its arrows, starting from the object that a displaced outsider loses. Writing that object as o^1, the cycle reads $o^1 \to i^1 \to o^2 \to \cdots \to o^m \to i^m \to o^1$, and the walk visits its vertices in this order. Agent i^m, who takes o^1, receives something he did not hold before and is therefore strictly better off, so at least one agent of the cycle strictly improves; consider the first such agent on the walk, that is, the one with the smallest index. Every agent met before him receives exactly what he already held, so each object o^{r+1} passed on the way is the current allotment of the preceding agent i^r. Each of those objects must be individually owned by the agent it points to: it is assigned in the reduced allocation to the preceding agent on the walk, and an assigned public object points to its holder, so an object pointing elsewhere cannot be public. This is the one place where the edge rule for assigned public objects in Definition 5.1 is essential. The walk therefore terminates at the first strictly improving agent and delivers the required ownership chain.

Lemma B.1 (Cycle chains) *Consider Step t of the $(\triangleright, \mu)$-STTC algorithm, and let $\mathcal{C}$ be an ordinary cycle forming at that step. Write $\mu^{\triangleright} := \mathrm{STTC}(R, \mu, \triangleright)$. Let μ' be the reduced allocation at Step t, which assigns each agent $i \in N^t$ his allotment $\mu(i)$ if $\mu(i) \in O^t \cup \{\varnothing\}$, and $\varnothing$ otherwise. Let ν be an allocation that agrees with μ' on N^t. Define the* ***improvers of $\mathcal{C}$ relative to ν*** *by $T(\mathcal{C}, \nu) := \{h \in S(\mathcal{C}) \mid \mu^{\triangleright}(h)\ P_h\ \nu(h)\}$. Suppose that there is an agent $j \in N^t \backslash S(\mathcal{C})$ such that $\nu(j)$ is an object of $\mathcal{C}$. Let m be the number of agents of $\mathcal{C}$ and set $o^1 := \nu(j)$. Enumerate $\mathcal{C}$ from that object as $o^1 \to i^1 \to o^2 \to i^2 \to \cdots \to o^m \to i^m \to o^1$. Then $T(\mathcal{C}, \nu) \neq \varnothing$ and, letting q be the smallest index with $i^q \in T(\mathcal{C}, \nu)$,*

(a) *for each $r \in \{1, \ldots, q\}$, $o^r \in \omega_{i^r}$; and*
(b) *for each $r \in \{1, \ldots, q-1\}$, $\nu(i^r) = \mu^{\triangleright}(i^r) = o^{r+1}$.*

Proof. **[Proof of Lemma B.1]** Fix Step t of the $(\triangleright, \mu)$-STTC algorithm, an ordinary cycle $\mathcal{C}$ forming at that step, an allocation ν that agrees with the reduced allocation

μ' on N^t, and an agent $j \in N^t \backslash S(\mathcal{C})$ such that $\nu(j)$ is an object of $\mathcal{C}$. We adopt the enumeration $o^1 \to i^1 \to \cdots \to o^m \to i^m \to o^1$ of the statement.

Each $h \in S(\mathcal{C})$ points to his most-preferred object in $O^t \cup \{\varnothing\}$ and receives it. Thus, $\mu^{\triangleright}(h)$ is this object. Because $\nu(h)$ is his reduced allotment, it also belongs to $O^t \cup \{\varnothing\}$. Hence, for each $h \in S(\mathcal{C})$,

$$\mu^{\triangleright}(h) \; R_h \; \nu(h). \tag{5.3}$$

By antisymmetry, every $h \in S(\mathcal{C}) \backslash T(\mathcal{C}, \nu)$ therefore satisfies $\mu^{\triangleright}(h) = \nu(h)$.

Agent i^m points to o^1 in $\mathcal{C}$, so $\mu^{\triangleright}(i^m) = o^1 = \nu(j)$. Since ν is an allocation, $j \neq i^m$, and o^1 is a real object, we have $\nu(i^m) \neq o^1 = \mu^{\triangleright}(i^m)$. Together with (5.3) and antisymmetry, $\mu^{\triangleright}(i^m) \; P_{i^m} \; \nu(i^m)$, so $i^m \in T(\mathcal{C}, \nu)$. Hence $T(\mathcal{C}, \nu) \neq \varnothing$ and q is well defined with $q \leqslant m$.

We prove part (b) first because it is used in the proof of part (a).

(b) Let $r < q$. Since q is the smallest index with $i^q \in T(\mathcal{C}, \nu)$, we have $i^r \notin T(\mathcal{C}, \nu)$, and therefore $\mu^{\triangleright}(i^r) = \nu(i^r)$. Since i^r points to o^{r+1} in $\mathcal{C}$, we also have $\mu^{\triangleright}(i^r) = o^{r+1}$.

(a) Agent j and the agents of $\mathcal{C}$ are all present at Step t, so $|N^t| \geqslant 2$ and, for each $h \in N^t$, by Definition 5.2, $\omega'_h = \omega_h \cap O^t \subseteq \omega_h$. Fix $r \in \{1, \ldots, q\}$. If $r = 1$, the reduced allocation assigns $o^1 = \nu(j)$ to agent j; if $r \geqslant 2$, then by (b), it assigns $o^r = \nu(i^{r-1})$ to agent i^{r-1}. In both cases o^r is assigned in the reduced allocation to an agent different from i^r. Were o^r public in the reduced instance, it would point to that agent rather than to i^r. Hence o^r is individually owned in the reduced instance by the agent it points to, i.e., $o^r \in \omega'_{i^r} \subseteq \omega_{i^r}$.

The next lemma establishes a fixed-point property of exclusion-core-stable allocations. It is used in the proof of Proposition 5.6 to recover each exclusion-core allocation as an STTC outcome from the null allocation, and in the proof of Theorem 5.3 to obtain the fixed-point characterization.

Lemma B.2 (Fixed-point property of exclusion-core-stable allocations) *Let R be a problem and let μ be an exclusion-core-stable allocation. Then, for every priority order $\triangleright$, we have* $\mathrm{STTC}(R, \mu, \triangleright) = \mu$.

Proof. **[Proof of Lemma B.2]** Let R be a problem and let μ be an exclusion-core-stable allocation. Fix a priority order $\triangleright$ and set $\mu^{\triangleright} := \mathrm{STTC}(R, \mu, \triangleright)$. Suppose, to the contrary, that $\mu^{\triangleright} \neq \mu$.

We first show that the algorithm executes at least one non-μ-cycle. Suppose that every executed cycle is a μ-cycle. Then each agent who leaves in a cycle receives his μ-allotment, and a real object is removed only by its μ-holder. An agent i who never leaves in a cycle therefore has $\mu(i) = \varnothing$ and receives $\varnothing$. So $\mu^{\triangleright} = \mu$, contrary to assumption.

Let $\mathcal{C}$ be an earliest non-μ-cycle formed at Step $t^\star$ in the STTC graph $\mathcal{G}^{t^\star}$. Every cycle executed before Step $t^\star$ is a μ-cycle, so as above a real object removed before Step $t^\star$ was taken by its μ-holder. Hence, for every $i \in N^{t^\star}$, $\mu(i) \in O^{t^\star} \cup \{\varnothing\}$.

We now construct a coalition $\widehat{S}$ and an allocation $\widehat{\sigma}$ and show that $\widehat{S}$ exclusion-blocks μ via $\widehat{\sigma}$. First, define the coalition $\widehat{S} := \{i \in S(\mathcal{C}) \mid \mu^{\triangleright}(i)\ P_i\ \mu(i)\}$ and show that $\widehat{S} \neq \varnothing$.

By the construction of $\mu^{\triangleright}$ via the STTC algorithm, we know that at Step $t^\star$, each agent in $S(\mathcal{C})$ points to, and receives, his most-preferred object in $O^{t^\star} \cup \{\varnothing\}$. Since $\mathcal{C}$ is a non-μ-cycle, there exists a cycle-agent $i \in S(\mathcal{C})$ with $\mu^{\triangleright}(i) \neq \mu(i)$. By the preceding observation, $\mu(i) \in O^{t^\star} \cup \{\varnothing\}$. Hence, agent i pointing at $\mu^{\triangleright}(i) \neq \mu(i)$ in the $(\triangleright, \mu)$-STTC algorithm implies that $\mu^{\triangleright}(i)\ P_i\ \mu(i)$. Thus, $\widehat{S} \neq \varnothing$.

Second, define the allocation $\widehat{\sigma}$ by

$$\widehat{\sigma}(i) := \begin{cases} \mu^{\triangleright}(i) & \text{if } i \in \widehat{S}, \\ \varnothing & \text{if } i \notin \widehat{S} \text{ and there exists a } k \in \widehat{S} \text{ such that } \mu(i) = \mu^{\triangleright}(k), \\ \mu(i) & \text{otherwise.} \end{cases}$$

We now show that coalition $\widehat{S}$ exclusion-blocks μ via $\widehat{\sigma}$. By the definition of $\widehat{S}$, for each $i \in \widehat{S}$, $\widehat{\sigma}(i)\ P_i\ \mu(i)$. By the preceding observation, for each cycle-agent $i \in S(\mathcal{C})$, $\mu(i) \in O^{t^\star} \cup \{\varnothing\}$. Hence, for each cycle-agent $i \in S(\mathcal{C})$, $\mu^{\triangleright}(i)\ R_i\ \mu(i)$.

If $\widehat{\sigma}$ Pareto dominates μ, no agent outside $\widehat{S}$ is worse off at $\widehat{\sigma}$. Therefore, the exclusion-enforcement condition is vacuously satisfied. Since every member of $\widehat{S}$ is strictly better off, $\widehat{S}$ exclusion-blocks μ via $\widehat{\sigma}$, establishing the desired contradiction.

If $\widehat{\sigma}$ does not Pareto dominate μ, let $j \notin \widehat{S}$ be any agent such that $\mu(j)\ P_j\ \widehat{\sigma}(j)$. Since every member of $\widehat{S}$ is strictly better off at $\widehat{\sigma}$, such an agent exists. By the definition of $\widehat{\sigma}$, $\widehat{\sigma}(j) = \varnothing$, and there exists an agent $i \in \widehat{S}$ such that $\mu(j) = \mu^{\triangleright}(i)$. In particular, $\mu(j)$ is a real object.

We first show that $j \notin S(\mathcal{C})$. Recall that every cycle agent $h \in S(\mathcal{C})$ satisfies $\mu^{\triangleright}(h)\ R_h\ \mu(h)$. If $j \in S(\mathcal{C}) \backslash \widehat{S}$, then strictness of preferences implies $\mu^{\triangleright}(j) = \mu(j)$. But $\mu^{\triangleright}(i) = \mu(j)$ for the distinct agent $i \in \widehat{S}$, contradicting feasibility of $\mu^{\triangleright}$. Hence, $j \notin S(\mathcal{C})$.

We next show that agent j is still present at Step $t^\star$. Suppose that j left at an earlier step. Since $\mathcal{C}$ is an earliest non-μ-cycle, agent j then received his μ-allotment, so $\mu^{\triangleright}(j) = \mu(j)$. Consequently, object $\mu(j)$ was removed before Step $t^\star$. But there exists a cycle agent $i \in \widehat{S}$ such that $\mu(j) = \mu^{\triangleright}(i)$. Since agent i receives this object in

$\mathcal{C}$ at Step $t^\star$, the object must still be available at that step, a contradiction. Therefore, $j \in N^{t^\star}$.

It remains to show that $\mu(j) \in \Omega(\widehat{S} \mid \mu)$.

Let μ' denote the reduced allocation at Step $t^\star$, which assigns each agent $h \in N^{t^\star}$ his allotment $\mu(h)$ if $\mu(h) \in O^{t^\star} \cup \{\varnothing\}$, and $\varnothing$ otherwise. Since $\mu(h) \in O^{t^\star} \cup \{\varnothing\}$ for every $h \in N^{t^\star}$, we have $\mu'(h) = \mu(h)$ for every such h. Hence, μ agrees with μ' on $N^{t^\star}$.

Moreover, there exists a cycle agent $i \in \widehat{S}$ such that $\mu(j) = \mu^{\triangleright}(i)$. Hence, $\mu(j)$ is a real object in $\mathcal{C}$; by the convention of Section 5.3.1, a cycle through the null object contains no real object, so $\mathcal{C}$ is ordinary. We have also shown that $j \in N^{t^\star} \backslash S(\mathcal{C})$ and that $T(\mathcal{C}, \mu) = \{h \in S(\mathcal{C}) \mid \mu^{\triangleright}(h)\ P_h\ \mu(h)\} = \widehat{S}$. Therefore, Lemma B.1 applies with $\nu := \mu$ and yields an enumeration $o^1 \to i^1 \to \cdots \to o^m \to i^m \to o^1$, where $o^1 = \mu(j)$, and an index $q \in \{1, \ldots, m\}$ such that $i^q \in \widehat{S}$ and, for each $r \leqslant q$, $o^r \in \omega_{i^r}$, and, for each $r < q$, $\mu(i^r) = o^{r+1}$. Therefore, by Lemma 5.5(d), applied with $S := \widehat{S}$, $\mu(j) = o^1 \in \Omega(\widehat{S} \mid \mu)$.

Since j was arbitrary among the agents outside $\widehat{S}$ who are worse off at $\widehat{\sigma}$, coalition $\widehat{S}$ exclusion-enforces $\widehat{\sigma}$ at μ. Since every member of $\widehat{S}$ is strictly better off, $\widehat{S}$ exclusion-blocks μ via $\widehat{\sigma}$, contradicting exclusion core stability of μ. Therefore, $\mu^{\triangleright} = \mu$.

Proof. [Proof of Proposition 5.6] Fix a problem R. By Proposition 5.5, the set of allocations $\{\mathrm{STTC}(R, \mu^{\varnothing}, \triangleright) \mid \triangleright \in \Delta\}$ is a subset of the exclusion core.

We now prove that the exclusion core is a subset of $\{\mathrm{STTC}(R, \mu^{\varnothing}, \triangleright) \mid \triangleright \in \Delta\}$. That is, we show that if an allocation μ belongs to the exclusion core, then there exists a priority order $\triangleright$ such that μ is the outcome of the $(\triangleright, \mu^{\varnothing})$-STTC algorithm.

Let μ be an allocation in the exclusion core. Consider an arbitrary priority order $\triangleright^*$. Without loss of generality, let $\triangleright^* : 1, 2, \ldots, n$. By Lemma B.2, $\mathrm{STTC}(R, \mu, \triangleright^*) = \mu$. Consequently, every cycle executed by the $(\triangleright^*, \mu)$-STTC algorithm is a μ-cycle. We use the steps of this algorithm to construct a priority order $\triangleright$ such that the $(\triangleright, \mu^{\varnothing})$-STTC algorithm also produces μ.

$(\triangleright^*, \mu)$-STTC algorithm:

Step 1. Starting with allocation μ, consider Step 1 of the $(\triangleright^*, \mu)$-STTC algorithm. Let $\mathscr{C}^1$ denote the set of all cycles that form, and define $S(\mathscr{C}^1) := \bigcup_{\mathcal{C} \in \mathscr{C}^1} S(\mathcal{C})$. Adjust the given priority order $\triangleright^*$ by moving the agents in $S(\mathscr{C}^1)$ above the remaining agents in $N^2 = N \backslash S(\mathscr{C}^1)$, preserving the relative order of agents in $S(\mathscr{C}^1)$ and in $N \backslash S(\mathscr{C}^1)$ according to $\triangleright^*$. Denote the adjusted priority order by $\triangleright^1$. We proceed inductively.

Step $t \geqslant 2$. Given the adjusted order $\triangleright^{t-1}$ at the end of Step $t-1$ of the $(\triangleright^*, \mu)$-STTC algorithm, consider Step t of the $(\triangleright^*, \mu)$-STTC algorithm. Let $\mathscr{C}^t$ denote the set of all cycles that form, and define $S(\mathscr{C}^t) := \bigcup_{\mathcal{C} \in \mathscr{C}^t} S(\mathcal{C})$. Adjust the current priority order

$\rhd^{t-1}$ by moving the agents in $S(\mathscr{C}^t)$ right below the agents in $S(\mathscr{C}^{t-1})$, and above the remaining agents $N^{t+1} = N \setminus \bigcup_{i=1}^{t} S(\mathscr{C}^i)$, preserving the relative order of agents in $S(\mathscr{C}^t)$ and in $N \setminus S(\mathscr{C}^t)$ according to $\rhd^{t-1}$. Denote the adjusted priority order by $\rhd^t$.

Finally, at the end of the $(\rhd^*, \mu)$-STTC algorithm, the resulting allocation is μ, and we obtain a final adjusted priority order $\rhd$. By construction of $\rhd$, if $t < s$, then every agent in $S(\mathscr{C}^t)$ has higher priority than every agent in $S(\mathscr{C}^s)$. Any agents who remain when the $(\rhd^*, \mu)$-STTC algorithm exhausts the set of real objects are ranked below all the sets $S(\mathscr{C}^t)$. Within each set, the relative order induced by $\rhd^*$ is preserved.

To complete the proof, starting with the null allocation $\mu^\varnothing$, we apply the $(\rhd, \mu^\varnothing)$-STTC algorithm and show that $\mathrm{STTC}(R, \mu^\varnothing, \rhd) = \mu$. Let T be the last step at which the $(\rhd^*, \mu)$-STTC algorithm executes a cycle. For each agent i, define $\tau(i)$ as follows. If there exists $t \in \{1, \ldots, T\}$ such that $i \in S(\mathscr{C}^t)$, let $\tau(i) := t$. If i receives the null object when the algorithm stops because no real object remains, let $\tau(i) := T + 1$. Thus, $\tau(i) < \tau(j)$ implies $i \rhd j$.

We prove by induction over the steps of the $(\rhd, \mu^\varnothing)$-STTC algorithm that every departing agent receives his μ-allotment. Consider Step t, and suppose that every agent who departed at an earlier step received his μ-allotment. We first show that the μ-allotment of every agent remaining at Step t is still available. Indeed, by the induction hypothesis a real object removed earlier was taken by its μ-holder, who has departed; so the μ-allotment of a remaining agent cannot have been removed. The null object is always available. Therefore, for every agent i remaining at Step t, $\mu(i) \in O^t \cup \{\varnothing\}$. Consider a cycle $\mathcal{C}$ in the current STTC graph.

Suppose first that exactly one agent i remains. Then, by Definition 5.2, $\omega_i' = O'$ and every remaining object is individually owned by agent i in the reduced instance, and no remaining object is public there. Reduced ownership therefore carries no information about original ownership, and we argue directly. Suppose that agent i points to an object $o \neq \mu(i)$. Since $\mu(i)$ is available and agent i points to his most-preferred object in $O^t \cup \{\varnothing\}$, $o \, P_i \, \mu(i)$. In the $(\rhd^*, \mu)$-STTC algorithm, agent i received $\mu(i)$ while the null object was available, so $\mu(i) \, R_i \, \varnothing$ and hence o is a real object. Since $o \, P_i \, \mu(i)$, object o cannot have been available when agent i received $\mu(i)$; it was therefore removed at an earlier step of that algorithm. As every cycle executed in that algorithm is a μ-cycle, there exists an agent j with $\tau(j) < \tau(i)$ such that $o = \mu(j)$. Since agent i is the only remaining agent, agent j has already departed, and by the induction hypothesis he departed with $\mu(j) = o$. Hence o is no longer available, a contradiction. Therefore agent i points to $\mu(i)$ and receives it.

Assume henceforth that at least two agents remain, so that, for each remaining agent k, $\omega_k' = \omega_k \cap O' \subseteq \omega_k$, and reduced private ownership implies original ownership.

First, suppose that $\mathcal{C}$ is not ordinary. By the convention of Section 5.3.1, $\mathcal{C}$ is then the self-cycle $i \to \varnothing \to i$ of a single agent i, who points to $\varnothing$. Since $\mu(i)$ is available and

the null object was also available when agent i was assigned in the $(\triangleright^*, \mu)$-STTC algorithm, strict preferences imply $\mu(i) = \varnothing$. Hence, agent i receives his μ-allotment.

Second, suppose that $\mathcal{C}$ is ordinary and contains an object that is public in the current reduced instance. Since the initial allocation is $\mu^{\varnothing}$, every reduced allocation assigns the null object to every remaining agent; hence no public object is assigned, and every public object points to the highest-priority remaining agent. Denote this agent by h. Then $\tau(h)$ is minimal among the τ-values of the remaining agents. We claim that h points to $\mu(h)$. Otherwise, h points to some object $o \neq \mu(h)$ with $o \, P_h \, \mu(h)$; since $\mathcal{C}$ is ordinary, o is a real object. Since h received $\mu(h)$ in the $(\triangleright^*, \mu)$-STTC algorithm, object o was removed before Step $\tau(h)$ of that algorithm. Hence, there exists an agent j with $\tau(j) < \tau(h)$ such that $o = \mu(j)$. By the induction hypothesis and the minimality of $\tau(h)$ among the remaining agents, agent j has already departed with object o, contradicting that o is still available. Therefore, h points to $\mu(h)$.

We now follow the cycle from h. If $\mu(h)$ is public in the current reduced problem, it points back to h. If $\mu(h)$ is individually owned by a remaining agent j, then $\mu(h)$ points to j. At Step $\tau(h)$ of the $(\triangleright^*, \mu)$-STTC algorithm, either j had already departed or j belonged to the same μ-cycle as h. The first possibility would imply $\tau(j) < \tau(h)$, contradicting the minimality of $\tau(h)$ among the remaining agents. Hence, $\tau(j) = \tau(h)$. The preceding argument then shows that j points to $\mu(j)$. Continuing around the cycle proves that every agent in $\mathcal{C}$ receives his μ-allotment.

Finally, suppose that $\mathcal{C}$ is ordinary and contains no object that is public in the current reduced instance. Thus, every object in $\mathcal{C}$ is individually owned by the agent to whom it points. Choose an agent $h \in S(\mathcal{C})$ with minimal $\tau(h)$ among the cycle agents. We again claim that h points to $\mu(h)$. Suppose instead that h points to an object $o \neq \mu(h)$. Since $\mu(h)$ is available, $o \, P_h \, \mu(h)$. Thus, object o was removed before Step $\tau(h)$ of the $(\triangleright^*, \mu)$-STTC algorithm. Let j be the owner to whom o points in the current cycle. When o was removed in the $(\triangleright^*, \mu)$-STTC algorithm, agent j had either departed at an earlier step or belonged to the cycle in which o was removed. In the first case, $\tau(j)$ is smaller than the step at which o was removed; in the second case, it equals that step. Since o was removed before Step $\tau(h)$, both cases imply $\tau(j) < \tau(h)$, contradicting the choice of h. Hence, h points to $\mu(h)$.

Since $\mu(h)$ belongs to $\mathcal{C}$ and no object in $\mathcal{C}$ is public in the current reduced instance, there exists a remaining agent j who individually owns $\mu(h)$. Hence, $\mu(h)$ points to j, and j belongs to $\mathcal{C}$. At Step $\tau(h)$ of the $(\triangleright^*, \mu)$-STTC algorithm, agent j had either already departed or belonged to the same μ-cycle as h. If j had already departed, then $\tau(j) < \tau(h)$, and since $j \in S(\mathcal{C})$, this contradicts the choice of h. Therefore, j belonged to the same μ-cycle as h, and $\tau(j) = \tau(h)$. Repeating the argument around the cycle shows that every cycle agent points to, and receives, his μ-allotment.

We have shown that at every step of the $(\triangleright, \mu^{\varnothing})$-STTC algorithm, all cycle agents receive their μ-allotments. The induction therefore continues until the algorithm termi-

nates. Suppose that no real object remains while some agents are still present. Suppose that a remaining agent i had a real μ-allotment. Since every removed object was taken by its μ-holder, object $\mu(i)$ could only have been taken by i himself, contradicting that i remains. Hence, every remaining agent has $\mu(i) = \varnothing$ and is assigned the null object. Consequently, $\mathrm{STTC}(R, \mu^{\varnothing}, \triangleright) = \mu$.

Appendix C: Proofs of the results in Section 5.4

Proof. [**Proof of Theorem 5.1**] Let ω be an endowment structure with a properly co-owned object $a \in O$. Thus, a has at least two co-owners, and there exists a non-co-owner $k \in N$. Let R be a preference profile in which each agent $i \in N$ satisfies $a \, P_i \, \varnothing \, P_i \, o$ for every object $o \in O \backslash \{a\}$; the relative ranking of the objects below $\varnothing$ is arbitrary. Let μ^a be the allocation at which object a is assigned to a non-co-owner k, while all other agents receive $\varnothing$; that is, $\mu^a(k) = a$ and, for each $i \in N \backslash \{k\}$, $\mu^a(i) = \varnothing$. The following two claims show that $\mu^a \in \mathrm{EC}(\omega, R)$ and $\mu^a \notin \mathrm{STTC}(\omega, R)$.

Claim C.1 $\mu^a \in \mathrm{EC}(\omega, R)$.

Proof. [**Proof**] Suppose, to the contrary, that a coalition $S \subseteq N$ exclusion-blocks μ^a via an allocation σ. Since each agent ranks only a above $\varnothing$, each member of S must receive a under σ in order to be strictly better off. In particular, the non-co-owner k, who already receives a at μ^a, cannot belong to S. Since some member of S receives a at σ, agent k is worse off at σ. Hence, exclusion-enforcement requires $a = \mu^a(k) \in \Omega(S \mid \mu^a)$. We show that this implies that every co-owner of a belongs to S. At μ^a, every co-owner of a receives $\varnothing$. The recursion defining $\Omega(S \mid \mu^a)$ adds an agent only when his μ^a-allotment belongs to the current direct eviction set, which is a subset of O and therefore never contains $\varnothing$. Consequently, a co-owner of a who does not belong to S cannot enter any of the recursively generated agent sets $S_k^{\mu^a}$. Since a belongs to the direct eviction set of an agent set only if that set contains all co-owners of a, the inclusion $a \in \Omega(S \mid \mu^a)$ therefore implies that every co-owner of a already belongs to S.

Since a is indivisible, at most one agent can receive a under σ. As a has at least two co-owners, not all members of S can be made strictly better off, contradicting that S exclusion-blocks μ^a. Therefore, $\mu^a \in \mathrm{EC}(\omega, R)$.

Claim C.2 $\mu^a \notin \mathrm{STTC}(\omega, R)$.

Proof. [**Proof**] Fix an arbitrary priority order $\triangleright$ and consider the first step of the $(\triangleright, \mu^{\varnothing})$-STTC algorithm. Since each agent ranks a above $\varnothing$, all agents initially point to a. Because a has at least two co-owners, at the first step it points to its highest-priority co-owner under $\triangleright$. Thus, a and its highest-priority co-owner form a self-cycle, which

assigns a to that co-owner. Once this self-cycle is executed, a is removed, so it cannot later be assigned to the non-co-owner k. Because $\triangleright$ was arbitrary, no priority order assigns a to k. Since $\mu^a(k) = a$, it follows that $\mu^a \notin \mathrm{STTC}(\omega, R)$.

The same argument applies to every owner-respecting flexible TTC algorithm described in Remark 5.9. Initially, all agents point to a, while a points to one of its co-owners. Thus, the first cycle containing a assigns it to a co-owner and removes it from the problem.

The same construction also shows that the fixed-point characterization in Theorem 5.3 fails under proper co-ownership. Initialize an owner-respecting flexible TTC algorithm at μ^a. All agents point to a. Because a is assigned at μ^a to the non-co-owner k, the owner-respecting rule makes a point to one of its co-owners. Hence, a and that co-owner form a self-cycle at the first step, so a is reassigned to the co-owner and removed. The resulting allocation differs from μ^a, although Claim C.1 establishes that μ^a is exclusion core stable.

By Remark 5.9, for every priority order $\triangleright$, $\mathrm{STTC}(\omega, R, \mu^{\varnothing}, \triangleright) \in \mathrm{EC}(\omega, R)$. Hence, $\mathrm{STTC}(\omega, R) \subseteq \mathrm{EC}(\omega, R)$. Thus, by Claims C.1 and C.2, $\mathrm{STTC}(\omega, R) \subsetneq \mathrm{EC}(\omega, R)$.

Proof. [**Proof of Corollary 5.2**] If the endowment structure is a mixed-ownership structure, then each associated problem R is a mixed-ownership problem and the coincidence follows from Proposition 5.6. Conversely, if the endowment structure is not a mixed-ownership structure, then some object $a \in O$ is properly co-owned; that is, at least two but not all agents own a. By Theorem 5.1, there then exists a preference profile R with $\mathrm{STTC}(\omega, R) \neq \mathrm{EC}(\omega, R)$.

Proof. [**Proof of Theorem 5.3**] Let R be a problem and let μ be an allocation. We prove (i) $\Rightarrow$ (iii) $\Rightarrow$ (ii) $\Rightarrow$ (i).

Assume (i). Then, for every priority order $\triangleright$, Lemma B.2 in Appendix B implies that $\mu = \mathrm{STTC}(R, \mu, \triangleright)$, which is (iii). That (iii) implies (ii) is immediate, since at least one priority order exists. Finally, assume (ii) and let $\triangleright$ be a priority order with $\mu = \mathrm{STTC}(R, \mu, \triangleright)$. Then Proposition 5.5 implies that μ is exclusion core stable, which is (i).

Proof. [**Proof of Lemma 5.9**] Let R be a problem, let μ and σ be allocations, and let S be a coalition that exclusion-blocks μ via σ. Since every member of S is strictly better off at σ, $S \subseteq S^+(\mu, \sigma)$. Consider an agent $j \notin S^+(\mu, \sigma)$ such that $\mu(j)\ P_j\ \sigma(j)$. Then, $j \notin S$. Since S exclusion-blocks μ via σ, $\mu(j) \in \Omega(S \mid \mu) \subseteq \Omega(S^+(\mu, \sigma) \mid \mu)$, where the inclusion follows from Lemma 5.5(b). Hence, $S^+(\mu, \sigma)$ exclusion-blocks μ via σ.

It remains to show that this blocking is self-contained. For ease of notation, let $S^+ := S^+(\mu, \sigma)$ and $\widehat{\sigma} := \widehat{\sigma}^{S^+, \mu}$. Each agent $i \in S^+$ receives $\widehat{\sigma}(i) = \sigma(i)$, which he strictly prefers to $\mu(i)$. Consider an agent $j \notin S^+$ such that $\mu(j)\ P_j\ \widehat{\sigma}(j)$. By the definition of the post-blocking allocation, $\widehat{\sigma}(j) = \varnothing$ and $\mu(j) \in \sigma(S^+) \backslash \{\varnothing\}$. Hence, there exists

an agent $i \in S^+$ such that $\sigma(i) = \mu(j)$. Since $j \notin S^+$, we have $i \neq j$; feasibility of σ therefore implies that $\sigma(j) \neq \mu(j)$. Since preferences are strict and $j \notin S^+$, it follows that $\mu(j) \; P_j \; \sigma(j)$. Because S^+ exclusion-blocks μ via σ, we obtain $\mu(j) \in \Omega(S^+ \mid \mu)$. Therefore, S^+ also exclusion-blocks μ via $\widehat{\sigma}$, and the blocking is self-contained.

Proof. [Proof of Theorems 5.6 and 5.11] Let R be a problem. Internal stability in Theorem 5.6 follows directly from the definition of the exclusion core. For the remaining claims, let μ be an allocation. If $\mu \in \mathrm{EC}(R)$, the constant sequence consisting only of μ is a self-contained exclusion-blocking path of length zero. Thus, for the path result, it remains to consider $\mu \notin \mathrm{EC}(R)$; this is also precisely the case relevant for external stability.

Fix a priority order $\rhd$ and apply the $(\rhd, \mu)$-STTC algorithm. Set $\mu^{\rhd} := \mathrm{STTC}(R, \mu, \rhd)$. By Proposition 5.5, $\mu^{\rhd}$ is exclusion core stable, and by Theorem 5.3, $\mu^{\rhd} \neq \mu$. If every executed cycle were a μ-cycle, every agent leaving in a cycle would receive his μ-allotment. Moreover, any agent remaining when no real object is left must have $\mu(i) = \varnothing$; otherwise, his real μ-allotment would have been removed in an earlier μ-cycle, in which he himself would have participated. Hence, the algorithm would return μ, contrary to $\mu^{\rhd} \neq \mu$. Thus, at least one non-μ-cycle forms. Let K be the number of steps at which at least one non-μ-cycle forms, and enumerate these steps as $t^1 < \cdots < t^K$. For each $\ell \in \{1, \ldots, K\}$, let $\mathscr{C}^\ell$ be the set of all non-μ-cycles at Step t^ℓ, and define $S(\mathscr{C}^\ell) := \bigcup_{\mathcal{C} \in \mathscr{C}^\ell} S(\mathcal{C})$.

Set $\mu^0 := \mu$. For each $\ell \in \{1, \ldots, K\}$, define μ^ℓ recursively by

$$\mu^\ell(i) := \begin{cases} \mu^{\rhd}(i), & \text{if } i \in S(\mathscr{C}^\ell), \\ \varnothing, & \text{if } i \notin S(\mathscr{C}^\ell) \text{ and there exists } h \in S(\mathscr{C}^\ell) \text{ such that } \mu^{\ell-1}(i) = \mu^{\rhd}(h), \\ \mu^{\ell-1}(i), & \text{otherwise.} \end{cases}$$

Call an agent outside $S(\mathscr{C}^\ell)$ an **outsider** at Step t^ℓ, and say that an outsider is **displaced** if some cycle agent receives his $\mu^{\ell-1}$-allotment. Since distinct cycles are disjoint in agents and in real objects, and every displaced outsider receives the null object, μ^ℓ is a well-defined allocation.

We use the same STTC steps to construct two kinds of coalitions. Define $\widehat{S}^0 := \varnothing$. For each $\ell \in \{1, \ldots, K\}$, define

- the *path coalition*

$$\widehat{S}^\ell_{\mathrm{path}} := \{i \in S(\mathscr{C}^\ell) \mid \mu^{\rhd}(i) \; P_i \; \mu^{\ell-1}(i)\},$$

  which compares the final STTC allotment with the current path allotment; and
- the *external-stability coalition*

$$\widehat{S}^\ell := \left\{ i \in \bigcup_{k=1}^{\ell} S(\mathscr{C}^k) \;\middle|\; \mu^{\rhd}(i) \; P_i \; \mu(i) \right\},$$

which compares the final STTC allotment with the original allotment.

Thus, the path coalition is local and myopic, whereas the external-stability coalition is cumulative and farsighted. By construction, for every $\ell \in \{1, \dots, K\}$, $\widehat{S}^{\ell-1} \subseteq \widehat{S}^{\ell}$.

History and location property. At each of the allocations $\mu^0, \dots, \mu^K$, every agent holds one of three allotments: his original allotment, his final STTC allotment, or the null object. More precisely, we record the following *history and location property*: for every $\ell \in \{1, \dots, K\}$ and every agent i,

(i) if i is present at Step t^ℓ, then $\mu^{\ell-1}(i) \in \{\mu(i), \varnothing\}$, and $\mu^{\ell-1}(i)$ is exactly his allotment in the reduced allocation at Step t^ℓ;
(ii) if i departed before Step t^ℓ, then $\mu^{\ell-1}(i)$ is his final STTC allotment, which is an already-removed object or the null object; and
(iii) if $i \in S(\mathscr{C}^\ell)$, then $\mu^{\triangleright}(i)\, R_i\, \mu^{\ell-1}(i)$.

To see this, note that once an agent participates in an STTC cycle he receives $\mu^{\triangleright}(i)$ and keeps it at every later μ^ℓ, whereas an agent who has not yet participated either keeps $\mu(i)$ or has been displaced and receives the null object. Part (iii) follows because a cycle agent points to his most-preferred object in $O^{t^\ell} \cup \{\varnothing\}$, and by part (i) his current allotment lies in that set.

Two consequences follow. By part (ii), any agent whose $\mu^{\ell-1}$-allotment is a real object still present at Step t^ℓ is himself present at Step t^ℓ. By part (iii) and antisymmetry, every cycle agent outside $\widehat{S}^{\ell}_{\mathrm{path}}$ keeps his current allotment and so displaces no outsider; hence the outsiders displaced in the definition of μ^ℓ are exactly those whose $\mu^{\ell-1}$-allotment is claimed by a member of $\widehat{S}^{\ell}_{\mathrm{path}}$. Hence, whenever $\widehat{S}^{\ell}_{\mathrm{path}} \neq \varnothing$, the post-blocking formula in Definition 5.7, applied to $\widehat{S}^{\ell}_{\mathrm{path}}$ at $\mu^{\ell-1}$ with blocking allocation μ^ℓ, reproduces μ^ℓ. If $\widehat{S}^{\ell}_{\mathrm{path}} = \varnothing$, then $\mu^\ell = \mu^{\ell-1}$.

We also use Lemma 5.5(c): once an agent's original allotment has been reached, the exclusion rights associated with his endowment become available to the coalition.

A cyclewise argument based on chains of dependence. Fix $\ell \in \{1, \dots, K\}$ and an ordinary cycle $\mathcal{C} \in \mathscr{C}^\ell$, and define its local improvers by

$$\widehat{S}^{\ell}_{\mathrm{path}}(\mathcal{C}) := \widehat{S}^{\ell}_{\mathrm{path}} \cap S(\mathcal{C}) = T(\mathcal{C}, \mu^{\ell-1}),$$

the last equality holding by the definitions of $\widehat{S}^{\ell}_{\mathrm{path}}$ and of $T(\mathcal{C}, \nu)$ in Lemma B.1. Suppose that there exists an agent $j \notin S(\mathcal{C})$ such that $\mu^{\ell-1}(j)$ is a real object of $\mathcal{C}$. Thus, j is displaced by this cycle. By parts (i) and (ii) of the history and location property, $j \in N^{t^\ell} \backslash S(\mathcal{C})$ and $\mu^{\ell-1}$ agrees with the reduced allocation on N^{t^ℓ}, so Lemma B.1 in Appendix B applies with $\nu := \mu^{\ell-1}$. By that lemma, $\widehat{S}^{\ell}_{\mathrm{path}}(\mathcal{C}) \neq \varnothing$. Let m be the number of agents of $\mathcal{C}$ and enumerate $\mathcal{C}$ from $o^1 := \mu^{\ell-1}(j)$ as $o^1 \to i^1 \to \cdots \to o^m \to i^m \to o^1$, and let q be the smallest index with $i^q \in \widehat{S}^{\ell}_{\mathrm{path}}(\mathcal{C})$. Then, for each $r \leqslant q$, $o^r \in \omega_{i^r}$, and, for each $r < q$, $\mu^{\ell-1}(i^r) = o^{r+1}$.

These are the conditions of a chain of dependence at $\mu^{\ell-1}$, and $\widehat{S}^{\ell}_{\mathrm{path}}(\mathcal{C})$ reaches i^q because i^q belongs to it. Therefore, by Lemma 5.5(d), applied with $S := \widehat{S}^{\ell}_{\mathrm{path}}(\mathcal{C})$ and the allocation $\mu^{\ell-1}$, $\mu^{\ell-1}(j) = o^1 \in \Omega(\widehat{S}^{\ell}_{\mathrm{path}}(\mathcal{C}) \mid \mu^{\ell-1})$. Since $\widehat{S}^{\ell}_{\mathrm{path}}(\mathcal{C}) \subseteq \widehat{S}^{\ell}_{\mathrm{path}}$, by Lemma 5.5(b), $\mu^{\ell-1}(j) \in \Omega(\widehat{S}^{\ell}_{\mathrm{path}} \mid \mu^{\ell-1})$.

Cycles involving the null object displace no outsider, so this argument is not needed for them. Since the ordinary cycles are disjoint, the preceding argument applies independently to every outsider displaced by the simultaneous cycles in $\mathscr{C}^{\ell}$. Therefore, whenever $\widehat{S}^{\ell}_{\mathrm{path}} \neq \varnothing$, coalition $\widehat{S}^{\ell}_{\mathrm{path}}$ exclusion-blocks $\mu^{\ell-1}$ via μ^{ℓ}. Together with the post-blocking formula above, this blocking is self-contained and establishes the myopic transition $\mu^{\ell-1} \mapsto \mu^{\ell}$. If $\widehat{S}^{\ell}_{\mathrm{path}} = \varnothing$, then $\mu^{\ell} = \mu^{\ell-1}$. When forming the final path, we delete this repeated allocation. This proves the path conclusion at every step; it remains to verify that the last path allocation is $\mu^{\triangleright}$, which we do below.

We now prove by induction that, for every $\ell \in \{1, \ldots, K\}$, $\widehat{S}^{\ell}$ exclusion-blocks μ via μ^{ℓ}. All cycles executed before Step t^1 are μ-cycles, so every agent who left earlier received his μ-allotment. Consequently, the μ-allotment of every agent present at Step t^1 is still available. Indeed, a real object removed earlier was taken by its μ-holder, who has departed, so the μ-allotment of a present agent cannot have been removed. Hence, every cycle agent at Step t^1 weakly prefers his final STTC allotment to his original allotment. Each cycle in $\mathscr{C}^1$ is a non-μ-cycle, so it contains an agent i with $\mu^{\triangleright}(i) \neq \mu(i)$; together with the preceding weak preference and antisymmetry, $\mu^{\triangleright}(i)\ P_i\ \mu(i)$. Since $\mu^0 = \mu$, we have $\widehat{S}^1 = \widehat{S}^1_{\mathrm{path}}$, and this set is nonempty. The cyclewise argument, now with $\mu^0 = \mu$, shows that every outsider harmed at μ^1 has his original allotment in $\Omega(\widehat{S}^1 \mid \mu)$. Thus, $\widehat{S}^1$ exclusion-blocks μ via μ^1.

Fix $\ell \in \{2, \ldots, K\}$, suppose that the claim holds at $\ell - 1$, and consider Step t^{ℓ}. Every member of $\widehat{S}^{\ell}$ has already participated in one of the cycles in $\mathscr{C}^1, \ldots, \mathscr{C}^{\ell}$ and therefore receives his final STTC allotment at μ^{ℓ}. By definition, he strictly prefers this allotment to his original allotment. It remains to verify that $\widehat{S}^{\ell}$ exclusion-enforces μ^{ℓ} at μ. Let $j \notin \widehat{S}^{\ell}$ satisfy $\mu(j)\ P_j\ \mu^{\ell}(j)$. If $\mu(j)\ P_j\ \mu^{\ell-1}(j)$, then the induction hypothesis and Lemma 5.5(b) give $\mu(j) \in \Omega(\widehat{S}^{\ell-1} \mid \mu) \subseteq \Omega(\widehat{S}^{\ell} \mid \mu)$.

Suppose instead that $\mu(j) P_j \mu^{\ell-1}(j)$ does not hold. Since preferences are strict, $\mu^{\ell-1}(j) R_j\ \mu(j)$. Together with $\mu(j)\ P_j\ \mu^{\ell}(j)$, we obtain $\mu^{\ell-1}(j)\ R_j\ \mu(j)\ P_j\ \mu^{\ell}(j)$. Thus, j becomes worse off relative to μ at the current step. Agent j cannot belong to $S(\mathscr{C}^{\ell})$. Indeed, suppose that $j \in S(\mathscr{C}^{\ell})$. Then j receives his final STTC allotment at μ^{ℓ}, and, by part (iii) of the history and location property, $\mu^{\ell}(j) = \mu^{\triangleright}(j)\ R_j\ \mu^{\ell-1}(j)$. But $\mu^{\ell-1}(j)\ R_j\ \mu(j)\ P_j\ \mu^{\ell}(j)$ implies $\mu^{\ell-1}(j)\ P_j\ \mu^{\ell}(j)$, which contradicts $\mu^{\ell}(j)\ R_j\ \mu^{\ell-1}(j)$. Since μ^{ℓ} differs from $\mu^{\ell-1}$ only for current cycle agents and displaced outsiders, j must be a newly displaced outsider. Hence, $\mu^{\ell}(j) = \varnothing$, and his real allotment $\mu^{\ell-1}(j)$ is taken by one of the current cycles. Since j is still present at Step t^{ℓ}, part (i) of the history and location property gives $\mu^{\ell-1}(j) \in \{\mu(j), \varnothing\}$. Moreover, $\mu^{\ell-1}(j)\ R_j\ \mu(j)\ P_j\ \mu^{\ell}(j) =$

$\varnothing$, which rules out $\mu^{\ell-1}(j) = \varnothing$. Hence, $\mu^{\ell-1}(j) = \mu(j)$. Let $\mathcal{C}$ be the cycle that displaces j. We show that $\mu(j) \in \Omega(\widehat{S}^{\ell} \mid \mu)$, first establishing that every local improver in $\mathcal{C}$ is either a member of $\widehat{S}^{\ell}$ or has his endowment contained in $\Omega(\widehat{S}^{\ell} \mid \mu)$, and then applying Lemmas B.1 and 5.5(d) to $\mathcal{C}$ with the ownership chain read relative to the original allocation μ.

For every local improver $i \in \widehat{S}^{\ell}_{\text{path}}(\mathcal{C})$, one of the following two cases applies. If $i \in \widehat{S}^{\ell}$, then i is already a member of the farsighted coalition. If $i \notin \widehat{S}^{\ell}$, then $\mu(i)\ R_i\ \mu^{\triangleright}(i)\ P_i\ \mu^{\ell-1}(i)$. Since i is still present at Step t^{ℓ}, part (i) of the history and location property gives $\mu^{\ell-1}(i) \in \{\mu(i), \varnothing\}$. The preference chain rules out $\mu^{\ell-1}(i) = \mu(i)$. Hence, $\mu^{\ell-1}(i) = \varnothing$, and therefore $\mu(i)\ P_i\ \mu^{\ell-1}(i)$.

Moreover, since i is still present at Step t^{ℓ}, he has not participated in any earlier cycle. Hence, $i \notin \widehat{S}^{\ell-1}$. By the induction hypothesis, $\widehat{S}^{\ell-1}$ exclusion-blocks μ via $\mu^{\ell-1}$. Since $i \notin \widehat{S}^{\ell-1}$ and $\mu(i)\ P_i\ \mu^{\ell-1}(i)$, the exclusion-enforcement condition implies $\mu(i) \in \Omega(\widehat{S}^{\ell-1} \mid \mu)$. Because $\widehat{S}^{\ell-1} \subseteq \widehat{S}^{\ell}$, Lemma 5.5(b) gives $\mu(i) \in \Omega(\widehat{S}^{\ell-1} \mid \mu) \subseteq \Omega(\widehat{S}^{\ell} \mid \mu)$. Lemma 5.5(c) then implies that the exclusion rights associated with the objects in ω_i are available to coalition $\widehat{S}^{\ell}$.

Return to the enumeration of $\mathcal{C}$ from $o^1 := \mu^{\ell-1}(j)$, $o^1 \to i^1 \to o^2 \to \cdots \to o^m \to i^m \to o^1$, supplied by Lemma B.1 with $\nu := \mu^{\ell-1}$, and let $q \in \{1, \ldots, m\}$ be the smallest index such that i^q is a local improver. Since $\mu^{\ell-1}(j) = \mu(j)$, as shown above, $o^1 = \mu(j)$. For every $r < q$, agent i^r is not a local improver. Hence, by Lemma B.1(b), $\mu^{\ell-1}(i^r) = o^{r+1}$. Since o^{r+1} is a real object, part (i) of the history and location property then implies $\mu(i^r) = \mu^{\ell-1}(i^r)$. Thus, for every $r < q$, $\mu(i^r) = \mu^{\ell-1}(i^r) = o^{r+1}$. By Lemma B.1(a), for each $r \leqslant q$, $o^r \in \omega_{i^r}$; and i^q is either a member of $\widehat{S}^{\ell}$ or satisfies $\omega_{i^q} \subseteq \Omega(\widehat{S}^{\ell} \mid \mu)$, as shown above. Therefore, by Lemma 5.5(d), applied with $S := \widehat{S}^{\ell}$ and the original allocation μ, $\mu(j) = o^1 \in \Omega(\widehat{S}^{\ell} \mid \mu)$.

This verifies exclusion-enforcement for every outsider harmed at μ^{ℓ}. Therefore, $\widehat{S}^{\ell}$ exclusion-blocks μ via μ^{ℓ}, completing the induction.

It remains to show that the final path allocation is the STTC allocation, $\mu^{K} = \mu^{\triangleright}$. Every agent who leaves in a μ-cycle retains his original allotment, which equals his final STTC allotment. Every agent who leaves in a non-μ-cycle receives his final STTC allotment at the corresponding path step or already holds it when that step leaves the path allocation unchanged. It remains only to consider an agent i who never leaves in a cycle and is assigned $\mu^{\triangleright}(i) = \varnothing$ when the algorithm stops because no real object remains. If $\mu(i) = \varnothing$, then, for every ℓ, $\mu^{\ell}(i) = \varnothing$. If $\mu(i)$ is a real object, then it was removed at some step and assigned to an agent h with $\mu^{\triangleright}(h) = \mu(i)$. The cycle of h cannot be a μ-cycle: otherwise $\mu(h) = \mu^{\triangleright}(h) = \mu(i)$, and since $\mu(i)$ is a real object and an allocation assigns a real object to at most one agent, this would imply $h = i$, contradicting that i never leaves in a cycle. Hence, for some ℓ', $h \in S(\mathscr{C}^{\ell'})$, and i is displaced at path step ℓ', so $\mu^{\ell'}(i) = \varnothing$. In both cases $\mu^{K}(i) = \varnothing = \mu^{\triangleright}(i)$. Consequently, $\mu^{K} = \mu^{\triangleright}$.

After deleting repeated allocations and relabeling the remaining sequence, the myopic transitions constructed above form a self-contained exclusion-blocking path from μ to the exclusion-core-stable allocation $\mu^{\triangleright}$. This proves Theorem 5.11.

Finally, applying the farsighted induction at $\ell = K$, $\widehat{S}^K$ exclusion-blocks the original allocation μ via μ^K. Since we have just shown that $\mu^K = \mu^{\triangleright}$, and $\mu^{\triangleright}$ belongs to the exclusion core, external stability follows. This proves Theorem 5.6.

Proof. [Proof of Corollary 5.12] Let R be a problem and μ an allocation. Since the sets of agents and objects are finite, the set of allocations is finite. If every allocation is exclusion core stable, the conclusion is immediate. Suppose, therefore, that at least one allocation lies outside the exclusion core. By Theorem 5.11, for each such allocation ν, choose one finite self-contained exclusion-blocking path from ν to an exclusion-core-stable allocation. Let L_ν denote its length and let $p_\nu > 0$ denote the probability, conditional on starting from ν, that the random process follows this path. Since the path is finite and every available transition is selected with positive probability, p_ν is positive. Since there are only finitely many allocations, $L := \max_{\nu \notin \mathrm{EC}(R)} L_\nu$ and $p := \min_{\nu \notin \mathrm{EC}(R)} p_\nu$ are well defined, and $p > 0$. Because every exclusion-core-stable allocation is absorbing, from every allocation outside the exclusion core the probability of reaching the exclusion core within the next L steps is at least p.

Conditional on the process not having reached the exclusion core after rL steps, its current allocation lies outside the exclusion core. The probability that it also fails to reach the exclusion core during the next L steps is therefore at most $1 - p$. Induction on r gives

$$\Pr(\text{the exclusion core has not been reached after } rL \text{ steps}) \leqslant (1-p)^r.$$

Since $(1-p)^r$ converges to zero, with probability one the process reaches an exclusion-core-stable allocation after finitely many steps and remains there.

Appendix D: An illustration of the STTC algorithm and stability results

The following example illustrates the STTC algorithm, the exclusion-core characterization in Proposition 5.6, the exclusion-core-stability test in Theorem 5.3, the external-stability conclusion of Theorem 5.6, and the path construction in Theorem 5.11.

Example D.1 (An illustration of the STTC algorithm and stability results) *Consider a set of agents $N = \{1,2,3,4,5,6,7\}$, a set of objects $O = \{a,b,c,d,e,f\}$, and endowments $\omega_1 = \{a,b\}$, $\omega_2 = \{c\}$, $\omega_3 = \{d\}$, and $\omega_4 = \omega_5 = \omega_6 = \omega_7 = \varnothing$. Thus, objects e and f are public. Fix the priority order $\triangleright : 1,2,3,4,5,6,7$ and the initial allocation $\boldsymbol{\mu} := (\boldsymbol{a},\boldsymbol{b},\boldsymbol{c},\boldsymbol{d},\boldsymbol{f},\boldsymbol{e},\boldsymbol{\varnothing})$. Let R be the preference profile displayed below, and define $\mu^{\triangleright} := \mathrm{STTC}(R,\mu,\triangleright)$. In the preference profile, boldface indicates each agent's allotment*

at μ, and underlining indicates his allotment at $\mu^{\triangleright}$. An allotment shown in both formats is unchanged by the algorithm.

$$\begin{array}{lccccc} R_1: & \underline{c} & b & \boldsymbol{a} & \varnothing & \dots \\ R_2: & \underline{\boldsymbol{b}} & \varnothing & \dots & & \\ R_3: & \boldsymbol{c} & \underline{d} & \varnothing & \dots & \\ R_4: & \boldsymbol{d} & \underline{a} & \varnothing & \dots & \\ R_5: & \underline{e} & d & \boldsymbol{f} & \varnothing & \dots \\ R_6: & \underline{f} & \boldsymbol{e} & \varnothing & \dots & \\ R_7: & a & \underline{\boldsymbol{\varnothing}} & \dots. & & \end{array}$$

Two kinds of superscripts are used below. A superscript t on N, O, ω, R, $\triangleright$, $\mathcal{G}$, $\mathscr{C}$, or $\tilde{\mu}$ refers to Step t of the algorithm. Thus, the primed notation N', O', ω', R', $\triangleright'$, and μ' from Definition 5.2 becomes N^t, O^t, ω^t, R^t, $\triangleright^t$, and $\tilde{\mu}^t$ here.

By contrast, a superscript ℓ on μ^ℓ indexes the allocations on the blocking path, with $\mu^0 = \mu$ and $\mu^K = \mu^{\triangleright}$. In this example, every step produces a non-μ-cycle, so the two indices happen to coincide; in general, they do not.

We now apply the STTC algorithm. At each step, we first describe the reduced STTC graph and the cycles that form. We then identify the corresponding myopic path coalition and cumulative external-stability coalition. At Step 1, we also illustrate the exclusion-core-stability test.

($\triangleright, \mu$)-STTC algorithm:

Step 1. *Let $N^1 = N$ and $O^1 = O$. Construct the STTC graph $\mathcal{G}^1 = \mathcal{G}(N^1, O^1, \omega, R, \mu, \triangleright)$ according to Definition 5.1. All agents point to their most-preferred object, which yields the edges $(1, c)$, $(2, b)$, $(3, c)$, $(4, d)$, $(5, e)$, $(6, f)$, $(7, a)$. Since agent 1 owns objects a and b, agent 2 owns object c, and agent 3 owns object d, edges $(a, 1)$, $(b, 1)$, $(c, 2)$, and $(d, 3)$ form. In addition, since public object f is assigned to agent 5 under μ, edge $(f, 5)$ forms, and since public object e is assigned to agent 6 under μ, edge $(e, 6)$ forms. Since no agent points to the null object $\varnothing$, it has no outgoing edge. The resulting directed STTC graph $\mathcal{G}^1$ is depicted in Figure 5.2.*

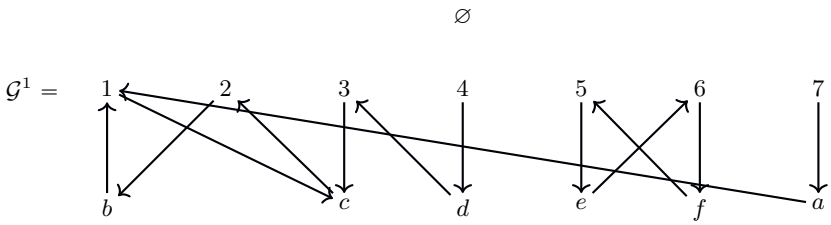

Figure 5.2: The STTC graph $\mathcal{G}^1$ constructed at Step 1.

In STTC graph $\mathcal{G}^1$, two cycles form: $1 \to c \to 2 \to b \to 1$ and $5 \to e \to 6 \to f \to 5$. These are earliest non-μ-cycles. We denote the set of all non-μ-cycles at this step by $\mathscr{C}^1$ and define

$$S(\mathscr{C}^1) := \bigcup_{\mathcal{C} \in \mathscr{C}^1} S(\mathcal{C}) = \{1, 2, 5, 6\}.$$

Execute these cycles by giving agent 1 object c, agent 2 object b, agent 5 object e, and agent 6 object f.

Exclusion-core-stability test. *The first non-μ-cycle $1 \to c \to 2 \to b \to 1$ identifies the following exclusion-blocking. Let $\widehat{S} = \{1\}$, and let the blocking allocation be $\widehat{\sigma} = (c, b, \varnothing, d, f, e, \varnothing)$.*

Agent $1 \in \widehat{S}$ receives c, while agent $3 \notin \widehat{S}$ (with $\mu(3) = c = \widehat{\sigma}(1)$) is left with $\varnothing$. Other agents receive their assignment at μ. To see that coalition $\widehat{S}$ exclusion-blocks μ via $\widehat{\sigma}$, notice that agent 1 is strictly better off receiving object c, which is assigned to agent 3 under μ. Agent 3, in turn, is worse off because his most-preferred object, c, is now assigned to agent 1. However, $b \in \omega_1$ and agent 2 holds b at μ; hence agent 2 becomes reachable from $\widehat{S} = \{1\}$. Since $c \in \omega_2$, it follows that $c \in \Omega(\widehat{S} \mid \mu)$. Thus, object c belongs to agent 1's extended eviction set. By Theorem 5.3, μ is not exclusion core stable.

We now construct the coalitions and allocations used for the path and external-stability results. Set $\mu^0 := \mu$.

Path transition. *Define the path coalition*

$$\widehat{S}^1_{\text{path}} = \{i \in S(\mathscr{C}^1) \mid \mu^{\triangleright}(i)\, P_i\, \mu^0(i)\} = \{1, 5, 6\}.$$

Note that agent $2 \notin \widehat{S}^1_{\text{path}}$ because $\mu^{\triangleright}(2) = b = \mu^0(2)$. Define allocation $\mu^1 = (c, b, \varnothing, d, e, f, \varnothing)$. Here, agents in $S(\mathscr{C}^1)$ receive their STTC allocations: $\mu^1(1) = \mu^{\triangleright}(1) = c$, $\mu^1(2) = \mu^{\triangleright}(2) = b$, $\mu^1(5) = \mu^{\triangleright}(5) = e$, and $\mu^1(6) = \mu^{\triangleright}(6) = f$. Agent 3 receives $\varnothing$ because $\mu^0(3) = c = \mu^{\triangleright}(1)$: his current allotment is taken by agent 1 in the cycle. All other agents keep their μ^0-allotments. Every member of $\widehat{S}^1_{\text{path}}$ is strictly better off at μ^1. Agent 3 is the only outsider who is worse off, and $\mu^0(3) = c \in \Omega(\widehat{S}^1_{\text{path}} \mid \mu^0)$, because c is owned by agent 2, who receives $b \in \omega_1$, and agent 1 belongs to $\widehat{S}^1_{\text{path}}$. Hence, $\widehat{S}^1_{\text{path}}$ exclusion-blocks μ^0 via μ^1. Applying Definition 5.7 to $\widehat{S}^1_{\text{path}}$ at μ^0 with blocking allocation μ^1 reproduces μ^1. Thus, the blocking is self-contained and establishes the transition $\mu^0 \mapsto \mu^1$.

External-stability coalition. *Define the external-stability coalition*

$$\widehat{S}^1 = \{i \in S(\mathscr{C}^1) \mid \mu^{\triangleright}(i)\, P_i\, \mu(i)\} = \{1, 5, 6\}.$$

At this step, $\widehat{S}^1 = \widehat{S}^1_{\text{path}}$ because both coalitions are defined relative to the original allocation $\mu = \mu^0$. Therefore, the preceding verification also shows that $\widehat{S}^1$ exclusion-blocks μ via μ^1.

Let $N^2 = N^1 \backslash \{1, 2, 5, 6\} = \{3, 4, 7\}$, $O^2 = O^1 \backslash \{b, c, e, f\} = \{a, d\}$.

***Step 2**. We derive the reduced problem for $N^2 = \{3, 4, 7\}$ and $O^2 = \{a, d\}$. The reduced endowment structure is given by $\omega^2_3 = \{d\}$ and $\omega^2_4 = \omega^2_7 = \varnothing$. Since $\omega_1 = \{a, b\}$, agent*

$1 \notin N^2$ and $a \in O^2$, object a turns into an unassigned public object in the reduced problem. The reduced preference profile is

$$\begin{array}{lllll} R_3^2: & d & \varnothing & & \\ R_4^2: & d & a & \varnothing & \\ R_7^2: & a & \varnothing. & & \end{array}$$

The reduced allocation, listed for agents 3, 4, and 7 in that order, is $\tilde{\mu}^2 = (\varnothing, d, \varnothing)$. The reduced priority order is $\rhd^2 : 3, 4, 7$.

Agents 3, 4, and 7 point to d, d, and a, respectively. Object d points to its owner, agent 3. Object a has no remaining owner and is unassigned in the reduced allocation, so it points to the highest-priority remaining agent, also agent 3. Since no agent points to the null object, it is isolated. The resulting graph is shown in Figure 5.3.

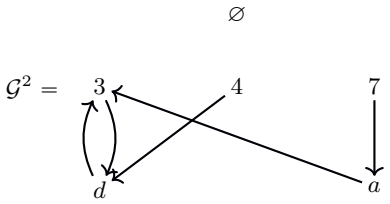

Figure 5.3: The STTC graph $\mathcal{G}^2$ constructed at Step 2.

In the STTC graph $\mathcal{G}^2$, only one cycle forms, $3 \to d \to 3$. Let $\mathscr{C}^2$ be the set containing this non-μ-cycle. Then $S(\mathscr{C}^2) = \{3\}$. We execute the cycle and assign object d to agent 3.

Path transition. *Define the path coalition*

$$\widehat{S}^2_{\text{path}} = \{i \in S(\mathscr{C}^2) \mid \mu^{\rhd}(i)\ P_i\ \mu^1(i)\} = \{3\}.$$

Agent 3 strictly improves from $\mu^1(3) = \varnothing$ to $\mu^{\rhd}(3) = d$. Define allocation $\mu^2 = (c, b, d, \varnothing, e, f, \varnothing)$. Here, agent $3 \in S(\mathscr{C}^2)$ receives $\mu^2(3) = \mu^{\rhd}(3) = d$. Agent 4 receives $\varnothing$ because $\mu^1(4) = d = \mu^{\rhd}(3)$: his current allotment is taken by agent 3. All other agents keep their μ^1-allotments. Agent 3 is strictly better off, and agent 4 is the only outsider who is worse off. Since $d \in \omega_3 \subseteq \Omega(\{3\} \mid \mu^1)$, coalition $\widehat{S}^2_{\text{path}} = \{3\}$ exclusion-blocks μ^1 via μ^2. Applying Definition 5.7 to $\widehat{S}^2_{\text{path}}$ at μ^1 with blocking allocation μ^2 reproduces μ^2. Thus, the blocking is self-contained and establishes the transition $\mu^1 \mapsto \mu^2$. Hence, we obtain the self-contained exclusion-blocking path $\mu \mapsto \mu^1 \mapsto \mu^2$, of length two, from μ to μ^2.

External-stability coalition. *Define the external-stability coalition*

$$\widehat{S}^2 = \{i \in \bigcup_{k=1}^{2} S(\mathscr{C}^k) \mid \mu^{\rhd}(i)\ P_i\ \mu(i)\} = \{1, 5, 6\}.$$

Note that $\widehat{S}^2 \neq \widehat{S}^2_{\text{path}}$, illustrating the key difference between path coalitions and external

stability coalitions. Specifically, agent $3 \notin \widehat{S}^2$ *because* $\mu(3) = c$, $\mu^{\triangleright}(3) = d$, *and he strictly prefers* c *to* d *(agent* 3 *is worse off at the final allocation than at* μ*). However,* $\widehat{S}^2$ *still exclusion-blocks* μ *via* μ^2*: all agents in* $\widehat{S}^2$ *are strictly better off, and agents* 3 *and* 4 *are worse off but their objects under* μ *are in the extended eviction set. Specifically,* $\mu(3) = c \in \Omega(\widehat{S}^2 \mid \mu)$*, as shown at Step 1. Therefore, by Lemma 5.5(c),* $\omega_3 = \{d\} \subseteq \Omega(\widehat{S}^2 \mid \mu)$*, and hence* $\mu(4) = d \in \Omega(\widehat{S}^2 \mid \mu)$.

Let $N^3 = N^2 \backslash \{3\} = \{4, 7\}$*, and* $O^3 = O^2 \backslash \{d\} = \{a\}$.

Step 3*. We derive the reduced problem for* $N^3 = \{4, 7\}$ *and* $O^3 = \{a\}$*. The reduced endowment structure is given by* $\omega_4^3 = \omega_7^3 = \varnothing$*. The reduced preference profile* R^3 *is:*

$$\begin{array}{lll} R_4^3: & a & \varnothing \\ R_7^3: & a & \varnothing. \end{array}$$

The reduced allocation, listed for agents 4 *and* 7 *in that order, is* $\tilde{\mu}^3 = (\varnothing, \varnothing)$*. The reduced priority order is* $\triangleright^3 : 4, 7$.

Agents 4 *and* 7 *both point to* a*. Since* a *is an unassigned public object, it points to the highest-priority remaining agent, agent* 4*. The null object is isolated. The resulting graph is shown in Figure 5.4.*

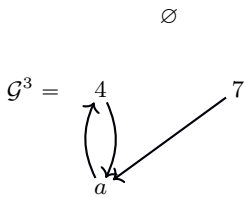

Figure 5.4: The STTC graph $\mathcal{G}^3$ constructed at Step 3.

In the STTC graph $\mathcal{G}^3$*, only one cycle forms,* $4 \to a \to 4$*. Let* $\mathscr{C}^3$ *be the set containing this non-*μ*-cycle. Then* $S(\mathscr{C}^3) = \{4\}$*. We execute the cycle and assign object* a *to agent* 4.

Path transition. *Define the path coalition*

$$\widehat{S}^3_{\text{path}} = \{i \in S(\mathscr{C}^3) \mid \mu^{\triangleright}(i) \; P_i \; \mu^2(i)\} = \{4\}.$$

Agent 4 *strictly improves from* $\mu^2(4) = \varnothing$ *to* $\mu^{\triangleright}(4) = a$*. Then, define allocation* $\mu^3 = (c, b, d, a, e, f, \varnothing) = \mu^{\triangleright}$*. Here, agent* $4 \in S(\mathscr{C}^3)$ *receives* $\mu^3(4) = \mu^{\triangleright}(4) = a$*. All other agents keep their* μ^2*-allotments. Agent* 4 *is strictly better off. Since object* a *is unassigned at* μ^2*, no outsider is worse off at* μ^3*, so the exclusion-enforcement condition is vacuously satisfied. Hence,* $\widehat{S}^3_{\text{path}} = \{4\}$ *exclusion-blocks* μ^2 *via* μ^3*. Applying Definition 5.7 to* $\widehat{S}^3_{\text{path}}$ *at* μ^2 *with blocking allocation* μ^3 *reproduces* μ^3*. Thus, the blocking is self-contained and establishes the transition* $\mu^2 \mapsto \mu^3$*. By Proposition 5.5,* $\mu^3 = \mu^{\triangleright}$ *is exclusion core stable. Therefore,*

$\mu = \mu^0 \mapsto \mu^1 \mapsto \mu^2 \mapsto \mu^3 = \mu^{\triangleright}$ is a self-contained exclusion-blocking path from μ to the exclusion core. Its path coalitions are $\widehat{S}^1_{\text{path}} = \{1, 5, 6\}$, $\widehat{S}^2_{\text{path}} = \{3\}$, and $\widehat{S}^3_{\text{path}} = \{4\}$. These are the myopic blocking coalitions associated with the three successive transitions.

External-stability coalition. *Define the external-stability coalition*

$$\widehat{S}^3 = \{i \in \bigcup_{k=1}^{3} S(\mathscr{C}^k) \mid \mu^{\triangleright}(i) \; P_i \; \mu(i)\} = \{1, 5, 6\} = \widehat{S}^2.$$

The external-stability coalition does not expand at Step 3. Agent 4 is a local improver because $\mu^2(4) = \varnothing$, $\mu^{\triangleright}(4) = a$, and he strictly prefers a to $\varnothing$; therefore, $4 \in \widehat{S}^3_{\text{path}}$. However, he is not a global improver because $\mu(4) = d$, $\mu^{\triangleright}(4) = a$, and he strictly prefers d to a; therefore, $4 \notin \widehat{S}^3$. This is precisely the distinction captured by the joint proof. Agent 4 was already displaced at Step 2, and $\mu(4) = d$ already belongs to $\Omega(\widehat{S}^2 \mid \mu)$. At Step 3, he improves locally by receiving the unassigned public object a, so no additional outsider is displaced and no additional exclusion rights are needed for the global blocking. Hence, the same coalition $\widehat{S}^3 = \widehat{S}^2 = \{1, 5, 6\}$ exclusion-blocks μ via the exclusion-core-stable allocation $\mu^3 = \mu^{\triangleright}$.

Termination. *Let $N^4 = N^3 \backslash \{4\} = \{7\}$ and $O^4 = O^3 \backslash \{a\} = \varnothing$. Since $N^4 \neq \varnothing$ and $O^4 = \varnothing$, we assign the null object to every agent in N^4, and the algorithm terminates after Step 3 with the final STTC allocation $\underline{\mu^{\triangleright}} = (\underline{c}, \underline{b}, \underline{d}, \underline{a}, \underline{e}, \underline{f}, \underline{\varnothing})$.*

Exclusion-core characterization. *We finally show that for problem R, the exclusion core consists of exactly two allocations: $\mu^{\triangleright}$, in which agent 4 receives a, and $\mu^{\star} = (c, b, d, \varnothing, e, f, a)$, in which agent 7 receives a. By Proposition 5.6, every exclusion-core-stable allocation can be obtained by running the STTC algorithm initialized at the null allocation $\mu^{\varnothing}$ with some priority order. Fix an arbitrary priority order $\triangleright'$. We now go through the $(\triangleright', \mu^{\varnothing})$-STTC algorithm and show that only two outcomes are possible.*

At Step 1 of the $(\triangleright', \mu^{\varnothing})$-STTC algorithm, agent 1 points to object c, object c points to its owner, agent 2, agent 2 points to object b, and object b points to its owner, agent 1. Hence, under every priority order, the non-$\mu^{\varnothing}$-cycle $1 \to c \to 2 \to b \to 1$ forms, assigning object c to agent 1 and object b to agent 2.

At Step 2, object c has been removed, so agent 3 points to his own object d and, under every priority order, forms the non-$\mu^{\varnothing}$-cycle $3 \to d \to 3$, receiving object d.

Next, consider the public objects e and f. As long as both objects remain, they point to the same agent, namely the highest-priority remaining agent. More generally, each remaining unassigned public object points to that agent. Only agent 5 points to e, and only agent 6 points to f. Hence, a cycle through e (resp. f) forms exactly when agent 5 (resp. agent 6) is the highest-priority remaining agent, in which case it is the self-cycle $5 \to e \to 5$ (resp. $6 \to f \to 6$). Since no object points to agent 5 or agent 6 before that moment, neither agent can leave the algorithm earlier. Consequently, under every priority order, agent 5 receives object e and agent 6 receives object f, at steps that depend on the priority order.

Finally, consider agents 4 *and* 7 *and object* a*. Once agent* 1 *has left at Step* 1*, object* a *has no remaining owner, becomes public in the reduced problem, and, being unassigned at* $\mu^{\varnothing}$*, points to the highest-priority remaining agent. Agent* 7 *points to object* a *from Step* 1 *on, and agent* 4 *points to object* a *from Step* 3 *on, once object* d *has been removed at Step* 2*. No other agent ever points to object* a*. Moreover, agents* 4 *and* 7 *point only to object* a *(or, for agent* 4*, to object* d *while it is present) or to the null object, so neither agent can receive any object other than* a*. Hence, object* a *is assigned in a self-cycle with whichever of agents* 4 *and* 7 *has the higher priority, once all higher-priority remaining agents have left, and the other agent receives* $\varnothing$ *when the algorithm terminates. In particular, if* $4 \rhd' 7$*, then agent* 4 *receives* a *and agent* 7 *receives* $\varnothing$*, yielding* $\mu^{\rhd}$*. If* $7 \rhd' 4$*, then agent* 7 *receives* a *and agent* 4 *receives* $\varnothing$*, yielding* $\mu^{\star}$*. Since one of these two cases applies to every priority order, the exclusion core consists of exactly* $\mu^{\rhd}$ *and* $\mu^{\star}$*.* $\diamond$

Appendix E: Acyclic endowment structures

Theorems 5.6 and 5.11 are established for mixed-ownership problems. This appendix examines how far they reach. We first characterize the acyclic endowment structures, a class that strictly contains the mixed-ownership structures and splits into two cases: the mixed-ownership structures themselves, and the structures under which a single group co-owns every object. The first case is covered by Theorems 5.6 and 5.11. For the second, we establish von Neumann–Morgenstern stability in Theorem E.5 and the existence of blocking paths in Theorem E.6. Corollaries E.7 and E.8 then combine the two cases. We then use two examples to show that acyclicity is not the boundary for external stability: outside this class, external stability may fail, but it may also hold at every preference profile. In both examples a self-contained exclusion-blocking path to the exclusion core nevertheless exists.

Acyclicity. Balbuzanov and Kotowski (2019) obtain positive results for relational economies whose priority structure is acyclic. Transcribing their condition into our setting yields the following restriction on endowment structures.

Definition E.1 (Acyclic endowment structures) *An endowment structure is* ***acyclic*** *if, for all distinct objects* $o, o' \in O$ *and all distinct agents* $i, j, k \in N$ *with* $i, k \in S^o$ *and* $j \notin S^o$*, we have* $k \in S^{o'}$ *and* $j \notin S^{o'}$*.*[14]

We thank Jacob Coreno for suggesting the following transparent description of acyclic endowment structures.

[14] Condition (1) of Balbuzanov and Kotowski (2019) states that a priority structure $\rhd$ is acyclic if for each $o' \neq o$, $i \rhd_o j$ and not $i \rhd_o k$ together imply $k \rhd_{o'} j$. Their Section 4.4 calls $\rhd$ a representation of an endowment structure if $i \rhd_o j$ exactly when $i \in S^o$ and $j \notin S^o$. Rewriting Condition (1) in terms of the sets S^o gives: if $i, k \in S^o$ and $j \notin S^o$, then $k \in S^{o'}$ and $j \notin S^{o'}$. Zhang (2026, Footnote 11) transcribes the same condition into the complex-endowment setting.

Proposition E.2 (A characterization of acyclic endowment structures) *An endowment structure is acyclic if and only if*

(a) *no object is properly co-owned, i.e., the structure is a mixed-ownership structure; or*
(b) *there is a coalition $T \subseteq N$ such that, for each object $o \in O$, $S^o = T$.*

Proof. [Proof of Proposition E.2] Fix an endowment structure ω. We first show that (a) implies acyclicity. Suppose that (a) holds and fix an object o. If o is public, then $S^o = N$, so no agent lies outside S^o. If o is individually owned by an agent h, then $S^o = \{h\}$, so no two distinct agents both lie in S^o. In either case there are no distinct agents i, j, k with $i, k \in S^o$ and $j \notin S^o$. Acyclicity requires something only of such agents, and o was arbitrary, so the structure is acyclic.

We next show that (b) implies acyclicity. Suppose that (b) holds with owning coalition T. Let o and o' be distinct objects and i, j, k distinct agents with $i, k \in S^o$ and $j \notin S^o$. Since $S^{o'} = T = S^o$, we have $k \in S^{o'}$ and $j \notin S^{o'}$. As the objects and agents were arbitrary, the structure is acyclic.

Conversely, suppose that the structure is acyclic; we show that (a) or (b) holds. If $|O| \leqslant 1$, then (b) holds: take $T := S^o$ for the unique object o when $|O| = 1$, and $T := N$ when $O = \varnothing$. Hence, suppose that $|O| \geqslant 2$ and that (a) fails, so that some object o satisfies $2 \leqslant |S^o| < |N|$.

Fix any object $o' \neq o$. To show $S^o \subseteq S^{o'}$, let $k \in S^o$ be arbitrary. Since $|S^o| \geqslant 2$, choose an agent $i \in S^o$ with $i \neq k$, and since $S^o \neq N$, choose an agent $j \notin S^o$. Then i, j, k are distinct, so acyclicity implies $k \in S^{o'}$. Hence, $S^o \subseteq S^{o'}$. To show the reverse inclusion, let $j \notin S^o$ be arbitrary. Since $|S^o| \geqslant 2$, choose distinct agents $i, k \in S^o$. Then acyclicity implies $j \notin S^{o'}$. Thus, every agent outside S^o also lies outside $S^{o'}$, which is equivalent to $S^{o'} \subseteq S^o$. Hence, $S^{o'} = S^o$. Since $o' \neq o$ was arbitrary, every object has the same set of owners, which is (b).

In case (b) with $|T| = 1$ a single agent owns every object, and with $T = N$ every object is public; both are mixed-ownership structures and are already covered by (a). The genuinely new case has $2 \leqslant |T| < |N|$ and, for each object $o \in O$, $S^o = T$; in this case one group co-owns everything and the remaining agents own nothing. We refer to it as **uniform co-ownership**. It is the situation of the "diarchy" of Balbuzanov and Kotowski (2019), where two co-monarchs jointly own all houses and the remaining agents do not own any real object.

Uniform co-ownership. Throughout the remainder of this paragraph the endowment structure is uniformly co-owned with owning coalition T, where $2 \leqslant |T| < |N|$; in particular, T is nonempty and may therefore serve as a blocking coalition. The key simplification is that exclusion rights are all-or-nothing.

Lemma E.3 (Collapse of the extended eviction sets) *Under uniform co-ownership, for each coalition S and each allocation μ,*

$$\Omega(S \mid \mu) = \begin{cases} O, & \text{if } T \subseteq S, \\ \varnothing, & \text{otherwise.} \end{cases}$$

Proof. [**Proof of Lemma E.3**] Let T be the owning coalition, so that $S^o = T$ for each object $o \in O$, and let S be a coalition and μ an allocation. By definition, $\omega(S) = \{o \in O \mid S^o \subseteq S\}$. If $T \nsubseteq S$, then $\omega(S) = \varnothing$, so $\mu^{-1}(\omega(S)) = \varnothing$ and $S_1^\mu = S$; the recursion is already at a fixed point and $\Omega(S \mid \mu) = \omega(S) = \varnothing$. If $T \subseteq S$, then, for each k, $S \subseteq S_k^\mu$, so $T \subseteq S_k^\mu$ and hence $\omega(S_k^\mu) = O$; therefore $\Omega(S \mid \mu) = O$.

For an allocation μ, let $A(\mu) := \{\nu \mid \text{for each } i \in T,\ \nu(i)\, P_i\, \mu(i)\}$ be the set of allocations at which every member of the owning coalition is strictly better off than at μ.

Lemma E.4 (Under uniform co-ownership) *Under uniform co-ownership, an allocation μ is exclusion core stable if and only if μ is Pareto efficient and $A(\mu) = \varnothing$.*

Proof. [**Proof of Lemma E.4**] Suppose that μ is Pareto efficient with $A(\mu) = \varnothing$, and suppose that some coalition S exclusion-blocks μ via σ. If $T \subseteq S$, then each member of T is strictly better off at σ, so $\sigma \in A(\mu)$, a contradiction. If $T \nsubseteq S$, then $\Omega(S \mid \mu) = \varnothing$ by Lemma E.3, so no agent outside S is worse off at σ; since $S \neq \varnothing$ and every member of S is strictly better off, σ Pareto dominates μ, contradicting Pareto efficiency.

Conversely, suppose that μ is exclusion core stable. If some allocation σ Pareto dominates μ, let $S := \{i \mid \sigma(i)\, P_i\, \mu(i)\}$, which is nonempty. No agent outside S is worse off at σ, so the exclusion-enforcement condition is vacuously satisfied and S exclusion-blocks μ via σ, a contradiction. Hence μ is Pareto efficient. If $\sigma \in A(\mu)$, define σ' such that for each $i \in T$, $\sigma'(i) = \sigma(i)$, and for each $j \notin T$, $\sigma'(j) = \varnothing$. This is an allocation, and each member of T is strictly better off at σ'. An agent $j \notin T$ is worse off only if $\mu(j)\, P_j\, \varnothing$, which requires $\mu(j) \in O = \Omega(T \mid \mu)$. Hence T exclusion-blocks μ via σ', a contradiction. Therefore $A(\mu) = \varnothing$.

Theorem E.5 (Von Neumann–Morgenstern stability under uniform co-ownership) *Under uniform co-ownership, the exclusion core is a von Neumann–Morgenstern stable set with respect to exclusion-domination.*

Proof. [**Proof of Theorem E.5**] Let R be a problem. Internal stability is immediate from the definition of the exclusion core: no allocation in $\mathrm{EC}(R)$ is exclusion-blocked, hence none is exclusion-dominated by any allocation.

For external stability, let $\mu \notin \mathrm{EC}(R)$. By Lemma E.4, either $A(\mu) \neq \varnothing$ or μ is not Pareto efficient. Suppose that μ is not Pareto efficient, and choose a Pareto efficient

allocation ν that Pareto dominates μ; such a ν exists because there are finitely many allocations. If $\nu \in \mathrm{EC}(R)$, then no agent outside $S := \{i \mid \nu(i)\ P_i\ \mu(i)\}$ is worse off at ν, so this nonempty coalition exclusion-blocks μ via ν; thus μ is exclusion-dominated by an allocation in $\mathrm{EC}(R)$, as required. If $\nu \notin \mathrm{EC}(R)$, then, since ν is Pareto efficient, Lemma E.4 gives $A(\nu) \neq \varnothing$; choose $\nu' \in A(\nu)$. For each $i \in T$, $\nu'(i)\ P_i\ \nu(i)\ R_i\ \mu(i)$, so $\nu' \in A(\mu)$. In every case that remains, $A(\mu) \neq \varnothing$.

We claim that $\mathrm{EC}(R) \cap A(\mu) \neq \varnothing$. Choose $\nu^0 \in A(\mu)$ and construct a sequence as follows. Given ν^k, stop if $\nu^k \in \mathrm{EC}(R)$. Otherwise, by Lemma E.4, either ν^k is not Pareto efficient, in which case let ν^{k+1} Pareto dominate ν^k; we call this a *step of the first kind*. Or $A(\nu^k) \neq \varnothing$, in which case let $\nu^{k+1} \in A(\nu^k)$; we call this a *step of the second kind*. In both cases, $\nu^{k+1} \in A(\mu)$: by induction, each $i \in T$ satisfies $\nu^k(i)\ P_i\ \mu(i)$; a step of the first kind gives $\nu^{k+1}(i)\ R_i\ \nu^k(i)$ and a step of the second kind gives $\nu^{k+1}(i)\ P_i\ \nu^k(i)$, so in either case $\nu^{k+1}(i)\ P_i\ \mu(i)$ by transitivity.

The construction terminates. Fix an agent $i_0 \in T$. Every step leaves i_0's allotment weakly better and every step of the second kind makes it strictly better, so there are at most $|O|$ steps of the second kind. These steps divide the sequence into at most $|O| + 1$ blocks, and every step within a block is of the first kind. At such a step the new allocation Pareto dominates the previous one, so within a block each allocation is Pareto dominated by the next. Since there are finitely many allocations and Pareto domination is transitive and irreflexive, each block is finite. As there are finitely many blocks, the construction terminates. Let ν^* be the terminal allocation. The construction stops only at an allocation in $\mathrm{EC}(R)$, so $\nu^* \in \mathrm{EC}(R)$; and every allocation in the sequence lies in $A(\mu)$, so $\nu^* \in A(\mu)$. Hence $\nu^* \in \mathrm{EC}(R) \cap A(\mu)$.

Finally, T exclusion-blocks μ via ν^*. Since $\nu^* \in A(\mu)$, each member of T is strictly better off. Let $j \notin T$ be worse off at ν^*. As ν^* is Pareto efficient, no agent is assigned an object ranked below the null object, for otherwise reassigning that agent to $\varnothing$ would be a Pareto improvement; hence $\nu^*(j)\ R_j\ \varnothing$ and therefore $\mu(j) \neq \varnothing$. Thus $\mu(j) \in O = \Omega(T \mid \mu)$ by Lemma E.3. Hence μ is exclusion-dominated by $\nu^* \in \mathrm{EC}(R)$.

Theorem E.6 (Paths to exclusion-core stability under uniform co-ownership) *Under uniform co-ownership, for each allocation μ there is a self-contained exclusion-blocking path from μ to the exclusion core.*

Proof. **[Proof of Theorem E.6]** Let R be a problem and let μ be an allocation. If $\mu \in \mathrm{EC}(R)$, the constant sequence consisting only of μ is a self-contained exclusion-blocking path of length zero. Hence, suppose that $\mu \notin \mathrm{EC}(R)$ and set $\mu^0 := \mu$. We construct a sequence recursively. Given μ^k, stop if $\mu^k \in \mathrm{EC}(R)$. Otherwise, proceed as follows.

First, suppose that μ^k is not Pareto efficient. Let μ^{k+1} Pareto dominate μ^k and set $S := \{i \mid \mu^{k+1}(i)\ P_i\ \mu^k(i)\}$, which is nonempty. No agent is worse off at μ^{k+1}, so

the exclusion-enforcement condition is vacuously satisfied and S exclusion-blocks μ^k via μ^{k+1}. Let $j \notin S$. Since μ^{k+1} Pareto dominates μ^k, we have $\mu^{k+1}(j)\ R_j\ \mu^k(j)$, and $j \notin S$ rules out strict preference, so $\mu^{k+1}(j) = \mu^k(j)$. Since $\mu^{k+1}(j) = \mu^k(j)$ and μ^{k+1} assigns each real object to at most one agent, $\mu^k(j) \notin \mu^{k+1}(S)\backslash\{\varnothing\}$. Therefore μ^{k+1} equals the post-blocking allocation induced by (S, μ^{k+1}) at μ^k, so the exclusion-blocking is self-contained. We call this a *step of the first kind*.

Second, suppose that μ^k is Pareto efficient. Since $\mu^k \notin \mathrm{EC}(R)$, Lemma E.4 implies that $A(\mu^k) \neq \varnothing$. Fix $\nu \in A(\mu^k)$ and define

$$\mu^{k+1}(i) := \begin{cases} \nu(i) & \text{if } i \in T, \\ \varnothing & \text{if } i \notin T \text{ and } \mu^k(i) \in \nu(T)\backslash\{\varnothing\}, \\ \mu^k(i) & \text{otherwise.} \end{cases}$$

This is an allocation: ν assigns each real object to at most one member of T, and an agent outside T retains $\mu^k(i)$ only if no member of T receives it. Each $i \in T$ satisfies $\mu^{k+1}(i) = \nu(i)\ P_i\ \mu^k(i)$. An agent $j \notin T$ is worse off only if $\mu^{k+1}(j) = \varnothing \neq \mu^k(j)$, which by Lemma E.3 occurs only if $\mu^k(j) \in \nu(T)\backslash\{\varnothing\} \subseteq O = \Omega(T \mid \mu^k)$. Hence T exclusion-blocks μ^k via μ^{k+1}. Since $\mu^{k+1}(T)\backslash\{\varnothing\} = \nu(T)\backslash\{\varnothing\}$, the displayed map is precisely the post-blocking allocation induced by (T, μ^{k+1}) at μ^k, so the exclusion-blocking is self-contained. Note also that $\mu^{k+1} \in A(\mu^k)$. We call this a *step of the second kind*.

Termination. Fix an agent $i_0 \in T$. Every step leaves i_0's allotment weakly better and every step of the second kind makes it strictly better, so there are at most $|O|$ steps of the second kind. These steps divide the sequence into at most $|O| + 1$ blocks, and every step within a block is of the first kind. At such a step the new allocation Pareto dominates the previous one, so within a block each allocation is Pareto dominated by the next. Since there are finitely many allocations and Pareto domination is transitive and irreflexive, each block is finite. As there are finitely many blocks, the construction terminates. By Lemma E.4, the terminal allocation is exclusion core stable. Since every transition in the constructed sequence is a self-contained exclusion-blocking, the sequence is a self-contained exclusion-blocking path from μ to the exclusion core.

Combining Proposition E.2 with Theorems 5.6 and E.5 yields the following.

Corollary E.7 (vNM stability under acyclic endowment structures) *Suppose that the endowment structure is acyclic. Then the exclusion core is a von Neumann–Morgenstern stable set with respect to exclusion-domination.*

Similarly, combining Proposition E.2 with Theorems 5.11 and E.6 yields the following.

Corollary E.8 (Paths to exclusion-core stability under acyclic endowment structures) *Suppose that the endowment structure is acyclic. Then, from each allocation, there is a self-contained exclusion-blocking path to the exclusion core.*

Corollary E.7 for the exclusion core of an object allocation problem with an acyclic endowment structure is the counterpart of Proposition 4 of Balbuzanov and Kotowski (2019), which establishes von Neumann–Morgenstern stability of the strong and weak exclusion cores of a relational economy with an acyclic priority structure. The same acyclicity condition supports both results, although the two solution concepts differ in general and coincide only under mixed ownership.

Acyclicity and the exclusion core characterization. Corollaries E.7 and E.8 do not weaken Corollary 5.2. Acyclicity is a domain for both conclusions but not for the characterization of the exclusion core by the STTC algorithms. Under uniform co-ownership with $2 \leqslant |T| < |N|$ every object is properly co-owned, so the construction proving Theorem 5.1 applies: it requires only a properly co-owned object and an agent who is not one of its co-owners, both of which are present. Mixed ownership therefore remains the maximal domain on which the STTC algorithms characterize the exclusion core, while it is not maximal for von Neumann–Morgenstern stability or for the existence of blocking paths.

Beyond acyclicity. Acyclicity is sufficient for the conclusions of Corollaries E.7 and E.8. Both examples below have non-acyclic endowment structures: in the first the exclusion core is not externally stable, in the second it is externally stable at every preference profile. Acyclicity is therefore not necessary for external stability. The examples also show that self-contained exclusion-blocking paths may exist under non-acyclic endowment structures; in both, such a path exists. Whether such a path always exists under proper co-ownership remains open. Although Lemma 5.7 is stated for mixed-ownership problems, its argument does not use the ownership structure: if σ Pareto dominates μ, no agent outside the coalition of agents who are strictly better off at σ is worse off, so that nonempty coalition exclusion-blocks μ via σ. Hence, in both examples below, every exclusion-core-stable allocation is Pareto efficient.

Example E.9 (Failure of external stability under a non-acyclic endowment structure) *Consider a problem with three agents $N = \{1, 2, 3\}$ and two objects $O = \{a, b\}$. Agent 1 owns object a, while agents 2 and 3 co-own object b; let ω denote this endowment structure. Since $2 \leqslant |S^b| < |N|$, object b is properly co-owned, so ω is not a mixed-ownership structure, and $S^a \neq S^b$, so the two objects do not share a common set of owners. By Proposition E.2, the endowment structure ω is therefore not acyclic. Let R be the following preference profile:*

$$\begin{array}{llll} R_1: & b & a & \varnothing \\ R_2: & a & b & \varnothing \\ R_3: & b & a & \varnothing \end{array}$$

Every agent ranks both real objects above the null object. Hence, if a real object were unassigned, it could be given to an agent receiving the null object, making that agent strictly better off without making anyone worse off. Therefore, every Pareto efficient allocation assigns both objects. There are six ways to assign a and b to two distinct agents. Two of them are Pareto dominated: $(a, b, \varnothing)$ by $(b, a, \varnothing)$, and $(\varnothing, b, a)$ by $(\varnothing, a, b)$. The remaining four are Pareto efficient:

1. *$\mu^1 = (b, a, \varnothing)$. At μ^1, agents 1 and 2 receive their most-preferred objects, so any blocking coalition must be $\{3\}$. Agent 3 can improve only by receiving a or b, both of which are assigned to agents outside the coalition. Since $\Omega(\{3\} \mid \mu^1) = \varnothing$, coalition $\{3\}$ has exclusion rights over neither object. Hence, μ^1 is not exclusion-blocked.*
2. *$\mu^2 = (a, \varnothing, b)$. At μ^2, agent 3 receives his most-preferred object and therefore cannot belong to a blocking coalition. Any blocking coalition must consequently be contained in $\{1, 2\}$. If it contains agent 1, then agent 1 must receive b to be strictly better off. However, for every $S \subseteq \{1, 2\}$, $b \notin \Omega(S \mid \mu^2)$, because agent 3, who co-owns b, does not belong to S. The only remaining possible blocking coalition is $\{2\}$, but $\Omega(\{2\} \mid \mu^2) = \varnothing$. Hence, μ^2 is not exclusion-blocked.*
3. *$\mu^3 = (\varnothing, a, b)$ is exclusion-blocked by $S = \{1\}$ via μ^2. Agent 1 is strictly better off. Agent 2 is the only outsider who is worse off: he loses object a. Since $a \in \Omega(\{1\} \mid \mu^3)$, coalition $\{1\}$ exclusion-enforces μ^2 at μ^3.*
4. *$\mu^4 = (b, \varnothing, a)$ is exclusion-blocked by $S = \{2, 3\}$ via μ^3. Both coalition members are strictly better off. Agent 1 is the only outsider who is worse off: he loses object b. Since $b \in \Omega(\{2, 3\} \mid \mu^4)$, coalition $\{2, 3\}$ exclusion-enforces μ^3 at μ^4.*

The preceding list contains all Pareto efficient allocations. Allocations μ^1 and μ^2 are not exclusion-blocked, whereas μ^3 and μ^4 are. Since every exclusion-core-stable allocation is Pareto efficient, it follows that $\mathrm{EC}(\omega, R) = \{\mu^1, \mu^2\}$. Allocation μ^4 lies outside the exclusion core, but no exclusion-core-stable allocation exclusion-dominates it. Relative to μ^4, only agent 2 is strictly better off at μ^1, so the only possible blocking coalition is $\{2\}$. Yet $\Omega(\{2\} \mid \mu^4) = \varnothing$, while agent 3 is worse off at μ^1. Thus, $\{2\}$ cannot exclusion-block μ^4 via μ^1. Similarly, only agent 3 is strictly better off at μ^2, but $\Omega(\{3\} \mid \mu^4) = \varnothing$, while agent 1 is worse off. Thus, $\{3\}$ cannot exclusion-block μ^4 via μ^2.

The exclusion core is therefore not externally stable. The two exclusion-blockings above are self-contained: in each case, the blocking allocation coincides with the induced post-blocking allocation. Hence, $\mu^4 \mapsto \mu^3 \mapsto \mu^2$ is a self-contained exclusion-blocking path of length two to the exclusion core. Thus, in this example, external stability fails even though μ^4 can reach the exclusion core through a self-contained exclusion-blocking path. $\diamond$

External stability may nevertheless hold outside the acyclic class, and may hold there at every preference profile, as the next example shows.

Example E.10 (External stability beyond acyclicity) *Consider three agents $N = \{1, 2, 3\}$ and two objects $O = \{a, b\}$. Agents 1 and 2 co-own object a, while object b is public; let ω*

denote this endowment structure. Since $2 \leqslant |S^a| < |N|$*, object* a *is properly co-owned, so* ω *is not a mixed-ownership structure, and* $S^a \neq S^b$*, so the two objects do not share a common set of owners. By Proposition E.2, the endowment structure* ω *is therefore not acyclic. We show that the exclusion core is nevertheless externally stable for* every *preference profile.*

Because $\omega(\{1,2\}) = \{a\}$*,* $\omega(N) = O$*, and, for every coalition* $S \notin \{\{1,2\}, N\}$*,* $\omega(S) = \varnothing$*, the extended eviction sets are*

$$\Omega(\{1,2\} \mid \mu) = \begin{cases} \{a,b\} & \text{if } \mu(3) = a, \\ \{a\} & \text{otherwise,} \end{cases} \qquad \Omega(N \mid \mu) = \{a,b\},$$

and, for every other coalition S*,* $\Omega(S \mid \mu) = \varnothing$*. Indeed, if* $\mu(3) = a$*, then agent 3 occupies an object in* $\omega(\{1,2\})$*, so the recursion adds him and* $\omega(\{1,2,3\}) = O$*.*

For a preference profile R *and an allocation* μ*, let* $A(\mu)$ *denote the set of allocations at which agents 1 and 2 are both strictly better off than at* μ*. The exclusion core admits the following description:*

$$\mu \in \mathrm{EC}(\omega, R) \text{ if and only if } \mu \text{ is Pareto efficient, and if } \mu(3) = a, \text{ then } A(\mu) = \varnothing. \quad (5.4)$$

To verify (5.4)*, suppose first that the right-hand side holds and that some coalition* S *exclusion-blocks* μ *via* σ*. If* $S = N$*, then every agent is strictly better off at* σ*, so* σ *Pareto dominates* μ*, contradicting Pareto efficiency. If* $\{1,2\} \nsubseteq S$*, then* $\Omega(S \mid \mu) = \varnothing$*, so no agent outside* S *is worse off at* σ*; as every member of* S *is strictly better off,* σ *again Pareto dominates* μ*. There remains* $S = \{1,2\}$*. If* $\mu(3) \neq a$*, then* $\Omega(\{1,2\} \mid \mu) = \{a\}$*, so agent 3 cannot be worse off and* σ *Pareto dominates* μ *once more. If* $\mu(3) = a$*, then agents 1 and 2 are both strictly better off at* σ*, so* $\sigma \in A(\mu)$*, contradicting* $A(\mu) = \varnothing$*.*

Conversely, suppose that $\mu \in \mathrm{EC}(\omega, R)$*. If some allocation* σ *Pareto dominates* μ*, then no agent outside* $S := \{i \mid \sigma(i)\ P_i\ \mu(i)\}$ *is worse off, so this nonempty coalition exclusion-blocks* μ *via* σ*; hence* μ *is Pareto efficient. Suppose that* $\mu(3) = a$ *and that* $\sigma \in A(\mu)$*. Define the allocation* $\widehat{\sigma}$ *by*

$$\widehat{\sigma}(i) := \begin{cases} \sigma(i) & \text{if } i \in \{1,2\}, \\ \varnothing & \text{if } i = 3 \text{ and } a \in \sigma(\{1,2\}), \\ a & \text{if } i = 3 \text{ and } a \notin \sigma(\{1,2\}). \end{cases}$$

Agents 1 and 2 are strictly better off at $\widehat{\sigma}$*, and agent 3 is worse off only when* $\widehat{\sigma}(3) = \varnothing \neq \mu(3)$*, in which case* $\mu(3) = a \in \Omega(\{1,2\} \mid \mu) = \{a,b\}$*. Thus* $\{1,2\}$ *exclusion-blocks* μ *via* $\widehat{\sigma}$*, contradicting* $\mu \in \mathrm{EC}(\omega, R)$*. Therefore no* $\sigma \in A(\mu)$ *exists, so* $A(\mu) = \varnothing$*.*

Given (5.4)*, external stability can be verified by exhaustive enumeration. There are thirteen allocations, and each agent has six strict preference orderings over* $O \cup \{\varnothing\}$*, yielding* $6^3 = 216$ *preference profiles. For each profile, we compute the exclusion core directly from Definition 5.4,*

confirm that it agrees with (5.4)*, and check that every allocation outside the exclusion core is exclusion-dominated by an allocation inside it. The verification succeeds for all 216 profiles.*[15]

The following profile illustrates the conclusion. Let R be given by

$$\begin{array}{llll} R_1: & a & b & \varnothing \\ R_2: & a & b & \varnothing \\ R_3: & a & b & \varnothing \end{array}$$

Here an allocation is Pareto efficient exactly when both objects are assigned, and among these it is exclusion core stable exactly when agent 3 does not hold a: if $\mu(3) = a$, then agents 1 and 2 hold b and $\varnothing$ in some order, and giving a to the one holding b and b to the one holding $\varnothing$ places an allocation in $A(\mu)$. Hence, by (5.4)*,*

$$\mathrm{EC}(\omega, R) = \big\{(a,b,\varnothing),\ (b,a,\varnothing),\ (a,\varnothing,b),\ (\varnothing,a,b)\big\},$$

and each of the nine remaining allocations is exclusion-dominated by a core allocation, as the following table records.

μ	S	ν	*agents worse off*	$\Omega(S \mid \mu)$
$(\varnothing,\varnothing,\varnothing)$	$\{1,2\}$	$(a,b,\varnothing)$	*none*	$\{a\}$
$(a,\varnothing,\varnothing)$	$\{2\}$	$(a,b,\varnothing)$	*none*	$\varnothing$
$(\varnothing,a,\varnothing)$	$\{1\}$	$(b,a,\varnothing)$	*none*	$\varnothing$
$(b,\varnothing,\varnothing)$	$\{2\}$	$(b,a,\varnothing)$	*none*	$\varnothing$
$(\varnothing,b,\varnothing)$	$\{1\}$	$(a,b,\varnothing)$	*none*	$\varnothing$
$(\varnothing,\varnothing,b)$	$\{1\}$	$(a,\varnothing,b)$	*none*	$\varnothing$
$(\varnothing,\varnothing,a)$	$\{1,2\}$	$(a,b,\varnothing)$	*agent 3*	$\{a,b\}$
$(b,\varnothing,a)$	$\{1,2\}$	$(a,b,\varnothing)$	*agent 3*	$\{a,b\}$
$(\varnothing,b,a)$	$\{1,2\}$	$(b,a,\varnothing)$	*agent 3*	$\{a,b\}$

In each of the first six rows every member of S is strictly better off at ν and no agent is worse off, so the exclusion-enforcement condition is vacuously satisfied. In each of the last three rows agents 1 and 2 are both strictly better off, and the only agent harmed is agent 3, whose allotment $\mu(3) = a$ lies in $\Omega(\{1,2\} \mid \mu) = \{a,b\}$. Moreover, in every row the blocking allocation ν coincides with the post-blocking allocation induced by (S,ν) at μ, so each exclusion-blocking is self-contained and each of the nine allocations reaches the exclusion core by a self-contained exclusion-blocking path of length one, whereas in Example E.9 external stability fails and μ^4 needs a path of length two. ◇

Together with Corollary E.7, Examples E.9 and E.10 show that acyclicity is sufficient but not necessary for external stability, and that non-acyclicity alone does not determine whether external stability holds. Example E.9 is also instructive for Theorem 5.11: there the exclusion core cannot be reached from μ^4 in one step, yet the

[15] The Python script used for this exhaustive verification is available from the authors upon request.

self-contained exclusion-blocking path $\mu^4 \mapsto \mu^3 \mapsto \mu^2$ does reach it. A self-contained exclusion-blocking path to the exclusion core may therefore exist where external stability fails.

Bibliography

Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., & Sönmez, T. (2005). The boston public school match. *American Economic Review*, *95*(2), 368–371.

Abdulkadiroğlu, A., & Sönmez, T. (1998). Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica: Journal of the Econometric Society*, *66*(3), 689–701.

Abdulkadiroğlu, A., & Sönmez, T. (1999). House allocation with existing tenants. *Journal of Economic Theory*, *88*(2), 233–260.

Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, *93*(3), 729–747.

Abebe, R., Cole, R., Gkatzelis, V., & Hartline, J. D. (2020). A truthful cardinal mechanism for one-sided matching. *Proceedings of the fourteenth annual ACM-SIAM symposium on discrete algorithms*, 2096–2113.

Akbarpour, M., Dworczak, P., & Kominers, S. D. (2024). Redistributive allocation mechanisms. *Journal of Political Economy*, *132*(6), 1831–1875. https://doi.org/10.1086/728111

Backhouse, R. E., & Medema, S. G. (2009). Retrospectives: On the definition of economics. *Journal of Economic Perspectives*, *23*(1), 221–33. https://doi.org/10.1257/jep.23.1.221

Balbuzanov, I., & Kotowski, M. H. (2019). Endowments, exclusion, and exchange. *Econometrica: Journal of the Econometric Society*, *87*(5), 1663–1692.

Balbuzanov, I., & Kotowski, M. H. (2024). The property rights theory of production networks. *Theoretical Economics*, *19*(4), 1619–1658. https://doi.org/10.3982/TE5413

Bentham, J. (1776). *A fragment on government* [Published anonymously]. T. Payne.

Bikhchandani, S., Chatterjee, K., & Sen, A. (2003). *Incentive compatibility in multi-unit auctions* (tech. rep.). UCLA, Decisions, Operations, and Technology Management. https://escholarship.org/uc/item/5rh9q995

Böhm-Bawerk, E. V. (2012). *The positive theory of capital*. Bubok.

Bosch-Domenech, A., & Sunder, S. (2000). Tracking the invisible hand: Convergence of double auctions to competitive equilibrium. *Computational Economics*, *16*(3), 257–284.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 17–33.

Colini-Baldeschi, R., de Keijzer, B., Leonardi, S., & Turchetta, S. (n.d.). Approximately efficient double auctions with strong budget balance. In *Proceedings of the 2016 annual acm-siam symposium on discrete algorithms (soda)* (pp. 1424–1443). https://doi.org/10.1137/1.9781611974331.ch98

Condorelli, D. (2013). Market and non-market mechanisms for the optimal allocation of scarce resources. *Games and Economic Behavior*, *82*, 582–591.

Correa, J. R., Figueroa, N., Lederman, R., & Stier-Moses, N. E. (2014). Pricing with markups in industries with increasing marginal costs. *Math. Program.*, *146*(1–2), 143–184. https://doi.org/10.1007/s10107-013-0682-8

Demange, G., Gale, D., & Sotomayor, M. (1986). Multi-item auctions. *Journal of Political Economy*, *94*(4), 863–872.

Deng, X., Goldberg, P., Tang, B., & Zhang, J. (2014). Revenue maximization in a bayesian double auction market. *Theoretical Computer Science*, *539*, 1–12. https://doi.org/https://doi.org/10.1016/j.tcs.2014.04.013

Deng, X., Sun, Y., Yin, M., & Zhou, Y. (2010). Mechanism design for multi-slot ads auction in sponsored search markets. *International Workshop on Frontiers in Algorithmics*, 11–22.

Deshmukh, K., Goldberg, A. V., Hartline, J. D., & Karlin, A. R. (2002). Truthful and competitive double auctions. *European Symposium on Algorithms*, 361–373.

Diamantoudi, E., Miyagawa, E., & Xue, L. (2004). Random paths to stability in the roommate problem. *Games and Economic Behavior*, *48*(1), 18–28. https://doi.org/10.1016/j.geb.2003.05.003

Dworczak, P., Kominers, S. D., & Akbarpour, M. (2021). Redistribution through markets. *Econometrica: Journal of the Econometric Society*, *89*(4), 1665–1698.

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, *97*(1), 242–259.

Ehlers, L. (2002). Coalitional strategy-proof house allocation. *Journal of Economic Theory*, *105*(2), 298–317.

Ehlers, L., Klaus, B., & Pápai, S. (2002). Strategy-proofness and population-monotonicity for house allocation problems. *Journal of Mathematical Economics*, *38*(3), 329–339. https://doi.org/https://doi.org/10.1016/S0304-4068(02)00059-9

Feldman, A. M. (1974). Recontracting stability. *Econometrica: Journal of the Econometric Society*, *42*(1), 35–44. Retrieved October 17, 2025, from http://www.jstor.org/stable/1913683

Finster, S., Loiseau, P., Mauras, S., Molina, M., & Pradelski, B. (2024). Equitable auctions. https://arxiv.org/abs/2403.07799

Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, *69*(1), 9–15.

Gallien, J. (2006). Dynamic mechanism design for online commerce. *Operations Research*, *54*(2), 291–310.

Gibbard, A. (1973). Manipulation of voting schemes: A general result. *Econometrica: Journal of the Econometric Society*, 587–601.

Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, *101*(1), 119–137.

Graafland, J. (2021). *Ethics and economics: An introduction to free markets, equality and happiness*. Routledge.

Groh, C.-C., & Reuter, M. (2023). *Mechanism design for unequal societies* (tech. rep.). ZEW Discussion Papers.

Groves, T. (1973). Incentives in teams. *Econometrica: Journal of the Econometric Society*, 617–631.

Haeringer, G., & Klijn, F. (2009). Constrained school choice. *Journal of Economic Theory*, *144*(5), 1921–1947.

Hafalir, I. E., Yenmez, M. B., & Yildirim, M. A. (2013). Effective affirmative action in school choice. *Theoretical Economics*, *8*(2), 325–363. https://doi.org/https://doi.org/10.3982/TE1135

Harsanyi, J. C. (1968). Games with incomplete information played by "bayesian" players, i-iii. part iii. the basic probability distribution of the game. *Management Science*, *14*(7), 486–502. Retrieved April 15, 2026, from http://www.jstor.org/stable/2628894

Hayek, F. A. (1945). The use of knowledge in society. *The American Economic Review*, *35*(4), 519–530. Retrieved April 15, 2026, from http://www.jstor.org/stable/1809376

Hurwicz, L. (1972). On informationally decentralized systems. *Decision and organization: A volume in Honor of J. Marschak*.

Hylland, A., & Zeckhauser, R. (1979). The efficient allocation of individuals to positions. *Journal of Political Economy*, *87*(2), 293–314.

Ishida, W., & Park, C. (2026). Group incentive-compatible allocation of discrete resources when ownership is partitioned. *Games and Economic Behavior*, *155*, 287–309. https://doi.org/10.1016/j.geb.2025.11.003

Klaus, B., & Klijn, F. (2007). Paths to stability for matching markets with couples. *Games and Economic Behavior*, *58*(1), 154–171. https://doi.org/10.1016/j.geb.2006.03.002

Kuhn, H. W. (1955). The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, *2*(1-2), 83–97.

Lavi, R., Mu'alem, A., & Nisan, N. (2003). Towards a characterization of truthful combinatorial auctions. *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, 574–583. https://doi.org/10.1109/SFCS.2003.1238230

Lehmann, D., Oćallaghan, L. I., & Shoham, Y. (2002). Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)*, *49*(5), 577–602.

Lin, P.-H., Brown, A. L., Imai, T., Wang, J. T.-y., Wang, S. W., & Camerer, C. F. (2020). Evidence of general economic principles of bargaining and trade from 2,000 classroom experiments. *Nature Human Behaviour*, *4*(9), 917–927.

Loertscher, S., & Mezzetti, C. (2019). The deficit on each trade in a vickrey double auction is at least as large as the walrasian price gap. *Journal of Mathematical Economics*, *84*, 101–106.

Mackenzie, A., & Zhou, Y. (2022). Tract housing, the core, and pendulum auctions. In *Graduate school of economics discussion paper series, no. e-22-005*. Kyoto University.

Mandler, M. (1999). *Dilemmas in economic theory: Persisting foundational problems of microeconomics*. Oxford University Press.

Maskin, E. S. (2008). Mechanism design: How to implement social goals. *American Economic Review*, *98*(3), 567–76. https://doi.org/10.1257/aer.98.3.567

McAfee, R. P. (1992). A dominant strategy double auction. *Journal of Economic Theory*, *56*(2), 434–450.

Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica: Journal of the Econometric Society*, *47*(1), 61–73. Retrieved April 15, 2026, from http://www.jstor.org/stable/1912346

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58–73.

Myerson, R. B., & Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, *29*(2), 265–281.

Nisan, N. (2015). Algorithmic mechanism design: Through the lens of multiunit auctions. In *Handbook of game theory with economic applications* (pp. 477–515, Vol. 4). Elsevier.

Nisan, N., & Ronen, A. (1999). Algorithmic mechanism design. *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, 129–140.

Núñez, M., & Rafels, C. (2005). The böhm–bawerk horse market: A cooperative analysis. *International Journal of Game Theory*, *33*(3), 421–430.

Peters, M., & Severinov, S. (2006). Internet auctions with many traders. *Journal of Economic Theory*, *130*(1), 220–245.

Rochet, J.-C. (1987). A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, *16*(2), 191–200.

Roth, A. E., & Postlewaite, A. (1977). Weak versus strong domination in a market with indivisible goods. *Journal of Mathematical Economics*, *4*(2), 131–137. https://doi.org/10.1016/0304-4068(77)90004-0

Roth, A. E., Sönmez, T., & Ünver, M. U. (2004). Kidney exchange. *The Quarterly Journal of Economics*, *119*(2), 457–488.

Roth, A. E., Sönmez, T., & Ünver, M. U. (2005a). Efficient kidney exchange: Coincidence of wants in a structured market.

Roth, A. E., Sönmez, T., & Ünver, M. U. (2005b). Pairwise kidney exchange. *Journal of Economic Theory*, *125*(2), 151–188.

Roth, A. E., & Vande Vate, J. H. (1990). Random paths to stability in two-sided matching. *Econometrica: Journal of the Econometric Society*, *58*(6), 1475–1480. https://doi.org/10.2307/2938326

Satterthwaite, M. A. (1975). Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, *10*(2), 187–217.

Say, J. B. (1836). *A treatise on political economy: Or the production, distribution, and consumption of wealth*. Grigg & Elliot.

Serrano, R., & Volij, O. (2008). Mistakes in cooperation: The stochastic stability of Edgeworth's recontracting. *Economic Journal*, *118*(532), 1719–1741.

Shapley, L., & Scarf, H. (1974). On cores and indivisibility. *Journal of Mathematical Economics*, *1*(1), 23–37. https://doi.org/10.1016/0304-4068(74)90033-0

Shapley, L. S., & Shubik, M. (1971). The assignment game i: The core. *International Journal of Game Theory*, *1*(1), 111–130.

Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, *70*(2), 111–137.

Sun, X., Tang, Q., Xiao, M., & Zhang, J. (2026). On the core of markets with co-ownerships and indivisibilities. *Journal of Economic Theory*, *232*, 106135. https://doi.org/10.1016/j.jet.2025.106135

Svensson, L.-G. (1999). Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, *16*(4), 557–567.

Triêu, A., Bos, I., Schröder, M., & Vermeulen, D. (2026). Maximal matchings. *Journal of Mathematical Economics*, *125*, 103254. https://doi.org/10.1016/j.jmateco.2026.103254

Varian, H. R. (2007). Position auctions. *international Journal of Industrial Organization*, *25*(6), 1163–1178.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, *16*(1), 8–37.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.

Zhang, J. (2026). *Consistency and the exclusion core in object allocation with co-ownership* [Working paper, arXiv:2511.21155]. https://arxiv.org/abs/2511.21155

Zhou, L. (1990). On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, *52*(1), 123–135.

Valorisation

Every day, scarce goods and services need to be allocated among people who want them: houses go to families, procurement contracts go to firms, school seats go to students, and donor kidneys go to patients. In many of these markets, the people involved know things the rest of the market does not: a buyer's true willingness to pay, a firm's true production cost, a student's true preference ranking. Left unaddressed, this asymmetry of information can cause markets to fail in ways that hurt people: a family that would happily pay for a house never gets the chance to, a hospital overpays for medical supplies because a supplier inflates its costs, or a student is matched to a school they did not actually prefer because the assignment procedure gave them a reason to misreport their true ranking. This thesis studies how the rules of such markets can be designed so that private information does not stand in the way of good outcomes from a societal perspective.

Relevance beyond academia. The markets studied in this thesis are not entirely abstract. They are markets nearly everyone interacts with over the course of their life. Consider the housing market. Whether a family renting or buying an apartment ends up matched with a suitable home depends on the rules governing how offers are solicited and accepted. When these rules are poorly designed, value-creating trades between buyers and sellers can fail to materialize simply because agents have an incentive to misrepresent their valuations or costs, or agents that place lower value on a property might get it over ones that value it more, while a policymaker's goal might be to serve as many families as possible, not just the ones who can pay the most. The mechanisms proposed in this thesis speak directly to this trade-off, providing implementable rules that a housing corporation, municipality, or online rental platform might want to adopt.

The same underlying problem shows up in procurement, where governments and firms award contracts to suppliers whose true production costs are private information. Poorly designed procurement auctions cost taxpayers and shareholders money, either through overpayment or through failing to award contracts to capable but overlooked suppliers. It shows up in the growing number of digital platforms that intermediate bilateral trade between many buyers and sellers, such as online marketplaces, freelance, and peer-to-peer electricity and energy-sharing schemes. In such a setting, a platform operator must decide how to match participants and set prices without

knowing participants' true costs or valuations, often under the added constraint that transactions must remain bilateral (a driver serves one passenger at a time, a solar panel owner sells excess electricity directly to a neighbor). The mechanisms proposed in this thesis are designed to be *implementable*: they can be adopted by a policymaker, platform operator, or market designer without requiring agents to volunteer information that they would otherwise have an incentive to hide.

The first part of the thesis concerns market design in the presence of a specific social priority: making sure that as many people as possible are served, rather than only maximizing the total value generated. This distinction matters most for essential goods and services, where availability and access, not just aggregate surplus, are the primary policy concern. Examples include housing for low- and middle-income households, medical supplies and equipment during periods of scarcity (such as during a public health crisis), and transportation services in underserved areas. A policymaker overseeing such a market may reasonably care more about how many households are housed, how many patients receive equipment, or how many residents have access to transportation, than about the total euro value of surplus generated.

The second part of the thesis, on markets with existing ownership structures, speaks to a different but equally common real-world problem: reallocating goods or positions that people already own, occupy, or have a claim to. Concrete examples include reassigning offices within a university or company as departments grow and shrink, reallocating school seats that become available mid-year due to transfers, house-swapping arrangements between existing tenants of housing corporations, and reallocating shared or inherited property among family members or co-owners. Because such reallocations must respect existing ownership or occupancy rights, naively "efficient" reallocation rules can be blocked or resisted by the very people they are meant to help, undermining the reallocation before it can even take place. The stability concept developed in this thesis, which builds on the idea of exclusion rights, helps identify reallocations that are not just efficient on paper but also robust to being unwound by the people involved, making them realistically implementable by an institution such as a university facilities office, a school district, or a housing corporation.

Target groups. The results of this thesis are primarily aimed at researchers in market design, mechanism design, and matching theory, who can build directly on the mechanisms, impossibility results, and algorithms developed here. Beyond academia, the results are relevant to a broader set of stakeholders. Policymakers and regulators overseeing markets for essential goods, such as housing corporations, healthcare procurement agencies, and municipal transport authorities, can draw on the mechanisms in this thesis when designing allocation rules that balance efficiency against access. Platform designers and operators of digital marketplaces, ride-hailing services, and peer-to-peer trading platforms can draw on the incentive-compatible mechanisms developed here to design matching and pricing rules that function well even when participants act strategically. Finally, institutions that periodically need to reallocate resources subject to existing ownership or occupancy rights, such as universities, school

districts, and housing corporations, can draw on the stability results from the second part of the thesis when designing internal reallocation procedures.

Activities and products. The theoretical mechanisms and algorithms developed in this thesis, including the Flip Algorithm proposed in Chapter 2 and the serial top trading cycles (STTC) procedure developed in Chapter 5, are, in principle, directly implementable as allocation rules or as the matching engine behind a digital platform. Because they are formally proven to be incentive compatible, they can be deployed without the need for extensive empirical testing of whether participants will attempt to manipulate the system, a property that is often costly and time-consuming to verify for ad hoc allocation rules designed without formal guarantees. This makes the results of this thesis directly usable as a starting point for practitioners designing or auditing real-world allocation systems, whether in a public agency, a housing corporation, or a technology platform.The results on mechanism subsidization may also give an indicator on possible social costs associated with implementation of such mechanisms.

Summary

Many economic interactions take place under private information: buyers know their own valuations, sellers know their own costs, and neither is directly observable to the other side of the market or to a social planner. When agents act strategically on this private information, mutually beneficial trades may fail to occur and desirable social outcomes may not be achieved. Mechanism design studies how the rules of interaction can be constructed so that agents' incentives align with a given social objective, despite this informational asymmetry.

This thesis applies mechanism design to matching markets with transfers, in which a policymaker or platform designer wishes to achieve a particular social goal, most commonly maximizing the number of value-creating trades, while ensuring that market participants are incentivized to truthfully reveal their private valuations and costs. The standard benchmark in the literature is the utilitarian, or total-utility-maximizing, rule, which selects the allocation that maximizes the aggregate surplus generated. While efficient in this aggregate sense, a total-utility maximizer can leave many value-creating trades unrealized, which is of particular concern when the goods or services being allocated are essential, such as housing, medical supplies, or transportation. The three chapters summarized below study, in a sequence of increasingly general market structures, whether and how implementable mechanisms can achieve more matches than total utility maximization, and what this costs in terms of forgone surplus.

Chapter 2: Maximal Matchings

The second chapter considers a market with multiple sellers, each offering one unit of a homogeneous, indivisible good, and multiple buyers, each demanding at most one unit. Both sellers' costs and buyers' valuations are private information. The social goal considered is to maximize the number of buyer–seller matches, referred to as a *maximal matching*, subject to the requirement that every individual trade be value-creating (*i.e.*, the buyer's valuation exceeds the seller's cost).

The chapter proposes a mechanism, the Flip Algorithm, with several desirable properties. First, the mechanism is incentive compatible: all agents are induced to truthfully report their valuations and costs. Second, among all maximal matchings, it selects the

one with the highest total utility. Third, matching maximization is shown to yield at most twice as many transactions as total utility maximization. Fourth, as with the well-known VCG mechanism for total utility maximization, the proposed mechanism is not budget balanced, meaning that a policymaker pursuing maximal matchings must be prepared to draw on outside funds. The chapter further shows that an equal split of surpluses under the proposed mechanism yields an outcome closely related to the egalitarian rule, connecting the matching-maximization objective to considerations of distributive justice.

Chapter 3: Monopolistic Trade Maximization

The third chapter turns to two related settings involving a single seller (monopolist): one in which the monopolist serves multiple buyers with unit demand, and one bilateral-monopoly setting in which the monopolist faces a single buyer with multi-unit demand. As in Chapter 2, both costs and valuations are private information, and the question is whether implementable, value-creating allocation rules exist that yield more transactions than total utility maximization.

When no assumption is placed on the shape of the seller's cost function (beyond positive marginal and total costs), the utilitarian rule turns out to be undominated in terms of the number of transactions: no implementable, value-creating rule can consistently improve on it. It is only when marginal cost is assumed to be non-decreasing in output that improvement becomes possible. The chapter then introduces the Average Allocation Rule (AAR) for this case, an incentive-compatible rule that dominates the utilitarian rule and is itself undominated in matching size. However, the AAR achieves this by permitting cross-subsidization across transactions, meaning it is not *bilaterally* value-creating (not every individual transaction generates a positive surplus on its own). The chapter shows that no implementable, bilaterally value-creating rule can always outperform total utility maximization in this setting. Finally, in the bilateral-monopoly case with multi-unit demand, total utility maximization and trade maximization are shown to coincide entirely. Taken together, these results indicate that, in contrast to the multiple-seller setting of Chapter 2, the utilitarian rule performs comparatively well in terms of trade volume when there is only a single seller.

Chapter 4: Heterogeneous Products

The fourth chapter generalizes the analysis to markets with heterogeneous, indivisible objects, in which agents may value different objects differently, moving beyond the homogeneous-good assumption of Chapters 2 and 3. This setting is closely related to the classical assignment-game model, though the chapter's focus is on static allocation rather than on dynamic price-adjustment processes.

Rather than proposing a specific mechanism, this chapter characterizes the broader class of implementable allocation rules in this multidimensional environment. Three natural economic properties are introduced: *monotonicity*, which ensures implementability; *value-creation*, which requires that objects are allocated only when doing so generates non-negative utility; and *irrelevance of unassigned objects*, which requires that lowering an agent's valuation for an object they do not receive does not affect the allocation. The chapter shows that any total-utility-maximizing rule satisfies all three properties, and establishes as its main result a cardinality restriction: any allocation rule satisfying these three properties that weakly dominates a maximal-cardinality total utility maximizer in the number of matches must in fact coincide with it in cardinality on every valuation profile.

Chapter 5: Exclusion Core Stability under Mixed Ownership

The final chapter turns from matching markets with transfers to object allocation problems without transfers, in which each agent already owns some initial endowment of indivisible objects and preferences over objects are ordinal rather than cardinal. Three classical benchmark models structure this literature: Shapley–Scarf housing problems, in which every agent owns exactly one object privately, house allocation problems, in which all objects are publicly owned, and house allocation with existing tenants problems lie between the two, combining privately owned objects, publicly owned objects, and agents who own nothing. This chapter studies a more general class, termed *mixed-ownership object allocation problems*, in which each object is either privately owned by a single agent or publicly owned by all agents, but proper co-ownership (an object jointly owned by a strict subset of agents) is excluded. This class nests all three benchmark models as special cases, while remaining tractable enough for a full characterization of stable allocations.

The relevant stability concept is the *exclusion core*, introduced by Balbuzanov and Kotowski (2019), which reinterprets ownership as a system of exclusion rights rather than as bundles of goods to be traded. Under this interpretation, a coalition can only block an allocation by reassigning objects it can access through owners' rights to exclude others, either directly or through chains of dependence. Unlike standard core concepts which can be empty (strong core) or admit wasteful allocations (weak core), the exclusion core is always nonempty and Pareto efficient, making it an appropriate stability concept even for more complex endowment structures such as mixed ownership.

To analyze exclusion core stability under mixed ownership, the chapter introduces the *serial top trading cycles (STTC) algorithm*, which, given a priority order over agents, an endowment structure, an initial allocation, and agents' preferences, constructs a sequence of directed graphs and executes top trading cycles sequentially. Unlike the generalized top trading cycles algorithm of Balbuzanov and Kotowski (2019), the STTC algorithm allows for an arbitrary initial allocation and relies on a single priority order

over agents rather than object-specific priority structures, which makes it well suited both to generating exclusion-core-stable allocations and to analyzing stability starting from an arbitrary allocation.

Building on two preliminary results that every STTC allocation is exclusion core stable, and that every exclusion-core-stable allocation can be reached by the STTC algorithm from the null allocation under some priority order, the chapter establishes four main results. First, once proper co-ownership is allowed, STTC allocations may form only a proper subset of the exclusion core, which implies that mixed ownership is the largest domain, among those defined by restrictions on endowment structures, on which the STTC algorithm fully characterizes the exclusion core. Second, the chapter provides a stability test: an allocation is exclusion core stable if and only if it is a fixed point of the STTC algorithm initialized at that allocation, regardless of the priority order used. Third, the exclusion core is shown to be a von Neumann–Morgenstern stable set. Fourth, the chapter establishes a dynamic result: from any allocation, there exists a self-contained exclusion-blocking path leading to the exclusion core, implying that a broad class of random blocking processes converges to an exclusion-core-stable allocation. The chapter further examines how far each of these results extends beyond mixed ownership, showing that exclusion core stability of STTC allocations holds for the full class of complex-endowment problems, while the characterization and fixed-point test fail as soon as a single object is properly co-owned. It also shows that von Neumann–Morgenstern external stability and the existence of blocking paths extend to the broader class of acyclic endowment structures, but not further.

Samenvatting

Veel economische interacties vinden plaats onder privé-informatie: kopers kennen hun eigen waarderingen, verkopers kennen hun eigen kosten, en geen van beide is direct waarneembaar voor de andere kant van de markt of voor een centrale planner. Wanneer partijen strategisch handelen op basis van deze privé-informatie, kunnen wederzijds voordelige transacties uitblijven en kunnen wenselijke maatschappelijke uitkomsten niet worden bereikt. Mechanism design bestudeert hoe regels en interactie zodanig kunnen worden vormgegeven dat de motivatie van partijen aansluiten bij een gegeven maatschappelijk doel, ondanks deze informatie-asymmetrie.

Dit proefschrift past mechanism design toe op matchingmarkten met geld, waarbij een beleidsmaker of platformontwerper een specifiek maatschappelijk doel wil bereiken, meestal het maximaliseren van het aantal waardecreërende transacties, terwijl marktpartijen tegelijkertijd worden gestimuleerd om hun private waarderingen en kosten naar waarheid te openbaren. De standaard maatstaf in de literatuur is de totale-nutmaximaliserende regel, die de allocatie kiest die het totale gegenereerde nut maximaliseert. Hoewel deze regel efficiënt is in de samengevoegde zin, kan een totale-nutmaximalisator veel waardecreërende transacties onbenut laten, wat met name zorgwekkend is wanneer de goederen of diensten essentieel zijn, zoals huisvesting, medische benodigdheden of vervoer. De drie hieronder samengevatte hoofdstukken onderzoeken, in een reeks van steeds algemenere marktstructuren, of en hoe implementeerbare mechanismen meer transacties kunnen bereiken dan totale-nutmaximalisatie en wat dit kost.

Hoofdstuk 2: Maximale Matchings

Het tweede hoofdstuk beschouwt een markt met meerdere verkopers, elk met één eenheid van een homogeen, ondeelbaar goed, en meerdere kopers, elk met een vraag naar ten hoogste één goed. Zowel de kosten van de verkopers als de waarderingen van de kopers zijn privé-informatie. Het maatschappelijke doel is het maximaliseren van het aantal transacties, aangeduid als een *maximale matching*, onder de voorwaarde dat elke individuele transactie waardecreërend is (*d.w.z.*, de waardering van de koper is groter dan de kosten van de verkoper).

Het hoofdstuk introduceert een mechanisme, het Flip-algoritme, met verschillende wenselijke eigenschappen. Ten eerste is het mechanisme implementeerbaar: alle partijen worden ertoe gestimuleerd hun waarderingen en kosten naar waarheid te rapporteren. Ten tweede selecteert het, onder alle maximale matchings, degene met het hoogste totale nut. Ten derde wordt aangetoond dat matchingmaximalisatie ten hoogste tweemaal zoveel transacties oplevert als totale-nutmaximalisatie, wat een bovengrens geeft aan de winst die behaald kan worden ten opzichte van de standaardmaatstaf. Ten vierde, net als bij het bekende VCG-mechanisme voor totale-nutmaximalisatie, is het voorgestelde mechanisme niet budgetneutraal, wat betekent dat een beleidsmaker die maximale matchings nastreeft, externe middelen moet kunnen aanwenden. Het hoofdstuk laat verder zien dat een gelijke verdeling van de surplussen onder het voorgestelde mechanisme een uitkomst oplevert die nauw verwant is aan de egalitaire regel, waarmee het doel van transactie maximalisatie wordt verbonden aan overwegingen van distributieve rechtvaardigheid.

Hoofdstuk 3: Monopolistische Handelsmaximalisatie

Het derde hoofdstuk richt zich op twee verwante settings met één verkoper (monopolist): één waarin de monopolist meerdere kopers met een vraag naar één eenheid bedient, en één bilaterale-monopolie-setting waarin de monopolist tegenover één koper met een vraag naar meerdere eenheden staat. Net als in hoofdstuk 2 zijn zowel kosten als waarderingen privé-informatie, en de vraag is of er implementeerbare, waardecreërende allocatieregels bestaan die meer transacties opleveren dan totale-nutmaximalisatie.

Wanneer geen aanname wordt gedaan over de vorm van de kostenfunctie van de verkoper (afgezien van positieve marginale kosten), blijkt de utilitaristische regel niet te domineren te zijn wat betreft het aantal transacties: geen enkele implementeerbare, waardecreërende regel kan hier consistent op verbeteren. Pas wanneer wordt aangenomen dat de marginale kosten niet-dalend zijn in de productieomvang, wordt verbetering mogelijk; het hoofdstuk introduceert hiervoor de Average Allocation Rule (AAR), een implementeerbare regel die de utilitaristische regel domineert en zelf niet gedomineerd wordt in matchinggrootte. De AAR bereikt dit echter door kruissubsidiëring tussen transacties toe te staan, wat betekent dat de regel niet *bilateraal* waardecreërend is (niet elke individuele transactie genereert op zichzelf een positief surplus). Het hoofdstuk laat zien dat geen enkele implementeerbare, bilateraal waardecreërende regel in deze setting altijd beter kan presteren dan totale-nutmaximalisatie. Tot slot blijkt in het geval van bilaterale monopolie met een vraag naar meerdere eenheden dat totale-nutmaximalisatie en handelsmaximalisatie volledig samenvallen. Deze resultaten laten gezamenlijk zien dat, in tegenstelling tot de setting met meerdere verkopers uit hoofdstuk 2, de utilitaristische regel relatief goed presteert in termen van handelsvolume wanneer er slechts één verkoper is.

Hoofdstuk 4: Heterogene Producten

Het vierde hoofdstuk generaliseert de analyse naar markten met heterogene, ondeelbare objecten, waarbij partijen verschillende objecten verschillend kunnen waarderen, en gaat daarmee verder dan de aanname van een homogeen goed uit hoofdstukken 2 en 3. Deze setting is nauw verwant aan het klassieke assignment-game-model, al ligt de focus van het hoofdstuk op statische allocatie in plaats van op dynamische prijsaanpassingsprocessen.

In plaats van een specifiek mechanisme voor te stellen, karakteriseert dit hoofdstuk de bredere klasse van implementeerbare allocatieregels in deze multidimensionale omgeving, waarin de private informatie van elke partij nu een vector van waarderingen over meerdere objecten is in plaats van een enkel getal. Drie natuurlijke economische eigenschappen worden geïntroduceerd: *monotoniciteit*, die implementeerbaarheid waarborgt; *waardecreatie*, die vereist dat objecten alleen worden toegewezen wanneer dit niet-negatief nut oplevert; en *irrelevantie van niet-toegewezen objecten*, die vereist dat het verlagen van de waardering van een partij voor een object dat zij niet ontvangt, de allocatie niet beïnvloedt. Het hoofdstuk laat zien dat elke totale-nutmaximaliserende regel aan alle drie deze eigenschappen voldoet, en stelt als hoofdresultaat een kardinaliteitsbeperking vast: elke allocatieregel die aan deze drie eigenschappen voldoet en een totale-nutmaximalisator van maximale kardinaliteit zwak domineert in het aantal matches, moet in feite met deze samenvallen in kardinaliteit voor elk waarderingsprofiel.

Hoofdstuk 5: Exclusion Core Stabiliteit bij Gemengd Eigendom

Het laatste hoofdstuk verlegt de aandacht van matchingmarkten met overdrachten naar objectallocatieproblemen zonder overdrachten, waarbij elke partij reeds een initieel bezit van ondeelbare objecten heeft en preferenties over objecten ordinaal in plaats van kardinaal zijn. Twee klassieke referentiemodellen structureren deze literatuur: Shapley–Scarf-huisvestingsproblemen, waarbij elke partij precies één object privé bezit, en huisvestingsallocatieproblemen, waarbij alle objecten publiek eigendom zijn. Huisvestingsallocatieproblemen met zittende huurders liggen tussen deze twee in, met een combinatie van privé-eigendom, publiek eigendom, en partijen die niets bezitten. Dit hoofdstuk bestudeert een algemenere klasse, aangeduid als *objectallocatieproblemen met gemengd eigendom*, waarbij elk object ofwel privé eigendom is van één partij, ofwel publiek eigendom is van alle partijen, maar waarbij eigenlijk mede-eigendom (een object dat gezamenlijk eigendom is van een strikte deelverzameling van partijen) is uitgesloten. Deze klasse omvat alle drie de bovengenoemde referentiemodellen als speciale gevallen, terwijl hij hanteerbaar genoeg blijft voor een volledige karakterisering van stabiele allocaties.

Het relevante stabiliteitsbegrip is de *exclusion core*, geïntroduceerd door Balbuzanov en Kotowski (2019), die eigendom herinterpreteert als een systeem van uitsluitingsrech-

ten in plaats van als bundels goederen die verhandeld kunnen worden. Onder deze interpretatie kan een coalitie een allocatie alleen blokkeren door objecten te herverdelen waartoe zij toegang heeft via de uitsluitingsrechten van eigenaren, hetzij direct, hetzij via ketens van afhankelijkheid. In tegenstelling tot de sterke core, die leeg kan zijn, en de zwakkere core, die inefficiënte allocaties kan bevatten zodra men de klassieke referentiemodellen verlaat, is de exclusion core altijd niet-leeg en Pareto-efficiënt, wat haar tot een geschikt stabiliteitsbegrip maakt, ook voor complexere eigendomsstructuren zoals gemengd eigendom.

Om exclusion core stabiliteit onder gemengd eigendom te analyseren, introduceert het hoofdstuk het *serial top trading cycles (STTC) algoritme*, dat, gegeven een prioriteitsvolgorde over partijen, een eigendomsstructuur, een initiële allocatie en de preferenties van partijen, een reeks gerichte grafen construeert en top trading cycles sequentieel uitvoert. In tegenstelling tot het generalized top trading cycles-algoritme van Balbuzanov en Kotowski (2019), staat het STTC-algoritme een willekeurige initiële allocatie toe en maakt het gebruik van één enkele prioriteitsvolgorde over partijen in plaats van objectspecifieke prioriteitsstructuren, waardoor het bijzonder geschikt is om zowel exclusion-core-stabiele allocaties te genereren als stabiliteit te analyseren vanuit een willekeurige allocatie.

Voortbouwend op twee voorbereidende resultaten, namelijk dat elke STTC-allocatie exclusion core stabiel is, en dat elke exclusion-core-stabiele allocatie kan worden bereikt door het STTC-algoritme vanuit de nulallocatie onder een geschikte prioriteitsvolgorde, stelt het hoofdstuk vier hoofdresultaten vast. Ten eerste, zodra er medeeigendom wordt toegestaan, kunnen STTC-allocaties slechts een echte deelverzameling van de exclusion core vormen, waaruit volgt dat gemengd eigendom, onder de domeinen gedefinieerd door beperkingen op eigendomsstructuren, het grootste domein is waarop het STTC-algoritme de exclusion core volledig karakteriseert. Ten tweede biedt het hoofdstuk een stabiliteitstoets: een allocatie is exclusion core stabiel dan en slechts dan als zij een fixpunt is van het STTC-algoritme geïnitialiseerd bij die allocatie, ongeacht de gebruikte prioriteitsvolgorde. Ten derde wordt aangetoond dat de exclusion core een von Neumann–Morgenstern stabiele verzameling is. Ten vierde stelt het hoofdstuk een dynamisch resultaat vast: vanuit elke allocatie bestaat er een self-contained exclusion-blocking pad dat naar de exclusion core leidt, wat impliceert dat een brede klasse van willekeurige blokkeringsprocessen convergeert naar een exclusion-core-stabiele allocatie. Het hoofdstuk onderzoekt vervolgens in hoeverre elk van deze resultaten geldig blijft buiten gemengd eigendom, en laat zien dat exclusion core stabiliteit van STTC-allocaties geldt voor de volledige klasse van problemen met complexe eigendomsstructuren, terwijl de karakterisering en de fixpunttoets al falen zodra één enkel object eigenlijk mede-eigendom is, en terwijl von Neumann–Morgenstern externe stabiliteit en het bestaan van blokkeringspaden zich wel uitstrekken tot de bredere klasse van acyclische eigendomsstructuren, maar niet verder.

About the author

Anh Triêu was born on August 27, 1999, in Vietnam. She obtained her BSc in Econometrics and Operations Research from Maastricht University's School of Business and Economics in 2021, followed by an MSc in Human Decision Science from the same institution in 2022. In September 2022, she began her PhD in Economics under the supervision of Dries Vermeulen, Iwan Bos and Marc Schröder, specializing in mechanism design and game theory.

Her research lies in microeconomic theory, with a focus on matching markets and mechanism design. She studies allocation problems under private information and institutional constraints, combining incentive compatibility with efficiency and stability criteria. Overall, her work aims to better understand how to design incentive-compatible, efficient, and stable mechanisms (or allocations) across different economic environments.

As part of her doctoral research, Anh visited Professor Bettina Klaus at the University of Lausanne, where she further developed her work on matching and market design. She has also presented her research at numerous international conferences and workshops, and since 2024, she has been part of the organizing team for the MLSE seminar in Maastricht.

In addition to her research, Anh has contributed to teaching as a tutor for several game-theory-related courses at Maastricht University.



ISBN 978-94-6534-722-6