Technische Universität München
TUM School of Management

Decision-Making in Competitive Environments

Computation, Selection and Inverse Optimization of Nash Equilibria in Mathematical Programming Games

Tobias Crönert, M.Sc.

Vollständiger Abdruck der von der TUM School of Management der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Wirtschafts- und Sozialwissenschaften (Dr. rer. pol.) genehmigten Dissertation.

Vorsitzender:		Prof. Dr. Maximilian Schiffer
Prüfer der Dissertation:	1.	Prof. Dr. Stefan Minner
	2.	Prof. Dr. Martin Bichler

Die Dissertation wurde am 16.08.2022 bei der Technischen Universität München eingereicht und durch die TUM School of Management am 15.11.2022 angenommen.

Acknowledgements

I would like to express my gratitude to everyone who supported me throughout my doctoral studies. First and foremost, I am truly grateful to my supervisor Prof. Dr. Stefan Minner for his support, guidance, availability, and valuable constructive feedback throughout my doctoral studies. I would like to express my gratitude to Prof. Dr. Martin Bichler for being a member of my examination committee, and Prof. Dr. Maximilian Schiffer for being the chairman of the committee.

Further, I would like to thank Prof. Christopher S. Tang, Ph.D., and Dr. Layla Martin for the joint work on one of the chapters of this thesis. I am also grateful to my collaborators Nicolas Busch, Moritz Rettinger, and Burakhan Sel – unfortunately, our joint project on data-driven decision making in commodity hedging is not part of this thesis.

I want to thank all current and former colleagues at the chair of Logistics and Supply Chain Management for their support, and in particular Dr. Patricia Rogetzer for her mentorship, and Sebastian Malicki for his help around the administration of the computing cluster.

I am grateful to Deutsche Forschungsgemeinschaft (DFG) for supporting my research as an associated member of the GRK 2201 ("Advanced Optimization in a Networked Economy"), and to all colleagues and advisors that are part of this research training group for the valuable discussions.

Finally, my deepest appreciation goes to my family for their support, advice, and encouragement throughout, but certainly not limited to my doctoral studies. Very special thanks go to Sina for her love, patience, and constant support.

Abstract

Competition is at the essence of most businesses and many managerial decisions include multiple competing decision makers. Popular examples include decisions on prices and production volumes, investments into capacities and technologies, or the location selection of production facilities, warehouses, or retail stores. We model such simultaneous, non-cooperative decision making in finite games. The predominant solution concept for a finite game is the identification of a Nash equilibrium. Prior algorithms are limited to normal-form games, or are restricted to the identification of a single equilibrium or all pure equilibria (which may fail to exist in general), but can neither enumerate all equilibria nor select the most likely equilibrium.

This thesis presents a solution method for finite games in a mathematical programming representation, i.e., action sets and payoffs are succinctly represented through inequality sets and objective functions, respectively. Our algorithm first determines all equilibria and subsequently identifies the most probable equilibrium according to the equilibrium selection theory by Harsanyi (1995). The algorithm is applied to a case study for hydrogen fuel station location planning in Munich, building on a novel Competitive flow capturing location-allocation model (C-FCLM). In contrast to prior formulations, this model requires no prior knowledge of competitor locations as competitors are expected to choose locations simultaneously in an emerging competitive environment. We show that fuel station providers who acknowledge the existence of their competitors and act accordingly can realize a profit increase of 17%.

We further present a novel inverse optimization approach for a subclass of finite games called Integer programming games (IPGs). This approach identifies parameter combinations that induce the observed (or desired) equilibrium solutions. We show that this inverse IPG corresponds to a bilevel problem which we solve using a cutting plane approach. Our approach extends prior methods for inverse optimization of integer programs to a competitive setting. We showcase its application for a competitive retail location selection problem: Whereas incumbent retailers can use ample historical data at a point of sales level to approximate customer attraction parameters and choose their

locations accordingly, new entrants lack this detailed information. They can, however, observe the resulting location structure of incumbents and deduct information on customer choice parameters using inverse optimization. We find that new entrants who base their location decision on inversely estimated parameters can improve their profits by 4-11% on average, compared with new entrants who rely on statistical averages for customer attraction parameters when making their location decision.

Contents

List of Tables

List of Figures

List of Abbreviations

BEV	Battery electric vehicle
C-FCLM	Competitive flow capturing location-allocation model
DFRLM	Deviation flow refueling location model
DM	Distribution mean
eSGM	exhaustive Sample generation method
EV	Electric vehicle
FCEV	Fuel cell electric vehicle
FCLM	Flow capturing location-allocation model
GNEP	Generalized Nash equilibrium problem
HD	Hamming distance
HFS	Hydrogen fuel station
IPG	Integer programming game
MCI	Multiplicative competitive interaction
MNL	Multinomial logit choice
MPG	Mathematical programmming game
NEP	Nash equilibrium problem
PoA	Price of Anarchy
PoE	Price of Equilibria
PoS	Price of Stability
SAA	Sample average approximation
SGM	Sample generation method
VCS	Value of the competitive solution
VSS	Value of the stochastic solution

Chapter 1

Introduction

1.1 Motivation

Economic decision making is hardly ever done in isolation, with companies and decision makers competing for talent, customers and market share, assets, raw material supply, or other resources. Yet, competition is often ignored in practice (Pang, 2010) and many optimization models neglect competitor actions or assume a static competitor decision, independent from the actions of the examined decision maker. It is apparent that a decision that is optimal under the pretense of a monopoly will hardly be optimal under competition. To derive a truly optimal decision, a company has to take the potential (re-)actions of competitors into account.

Within our work we model competitive decision making in finite games, i.e., games with a finite number of players (competitors) having a finite number of strategies at their disposal. For example, competitors might choose (integer) production volumes, where their capacities act as a natural upper bound. For such finite games, Nash (1951) proposes the solution concept of a stable fixpoint (later called the Nash equilibrium) based on three core ideas: A Nash equilibrium is a combination of strategies in which no player benefits from unilateral deviation (i). Any finite game has at least one Nash equilibrium in pure or mixed strategies (ii). In a non-cooperative, full information game between rational players with a single equilibrium, the unique equilibrium is the expected outcome of the game (iii).

This clear definition (i), general proof of existence (ii) and practical interpretation (iii) explain the continued popularity of Nash's solution concept (see, e.g., Nisan et al., 2007), yet, two main problems remain unsolved: An equilibrium can only be the expected outcome of the game if it is identifiable by all players (a). An equilibrium can only be the expected outcome of the game if it is unique (b). There have been many contributions partially addressing these two problems in the past, moving from simple zero-sum two-player games to increasingly realistic and complex settings. We review contributions to the identification of Nash equilibria (a) and equilibrium selection approaches (b) in detail in Section 2.1.

Within our work, we focus on a class of games for which neither of the two problems has been adequately addressed in prior work: Finite games in a mathematical programming representation (finite mathematical programming games). In contrast to (normal form) Nash games, in a Mathematical programmming game (MPG) potential strategies, strategy combinations or payoffs are not enumerated, but represented through a set of inequalities and objective functions, respectively. As such, a MPG resembles a Nash equilibrium problem (NEP), yet NEPs usually assume continuous decision variables whereas a MPG allows for discrete or integer decisions (Dragotto et al., 2021). Integer programming games (IPGs) (as introduced by Köppe et al., 2011) represent a special class of finite MPGs, in which all decision variables are integer (Köppe et al., 2011). In contrast to Generalized Nash equilibrium problems (GNEPs), in MPGs and NEPs the set of strategies per player does not depend on the decisions of other players. Figure 1.1 provides a graphic representation of this terminology. In the following, the term Nash equilibrium and equilibrium will be used interchangeably.

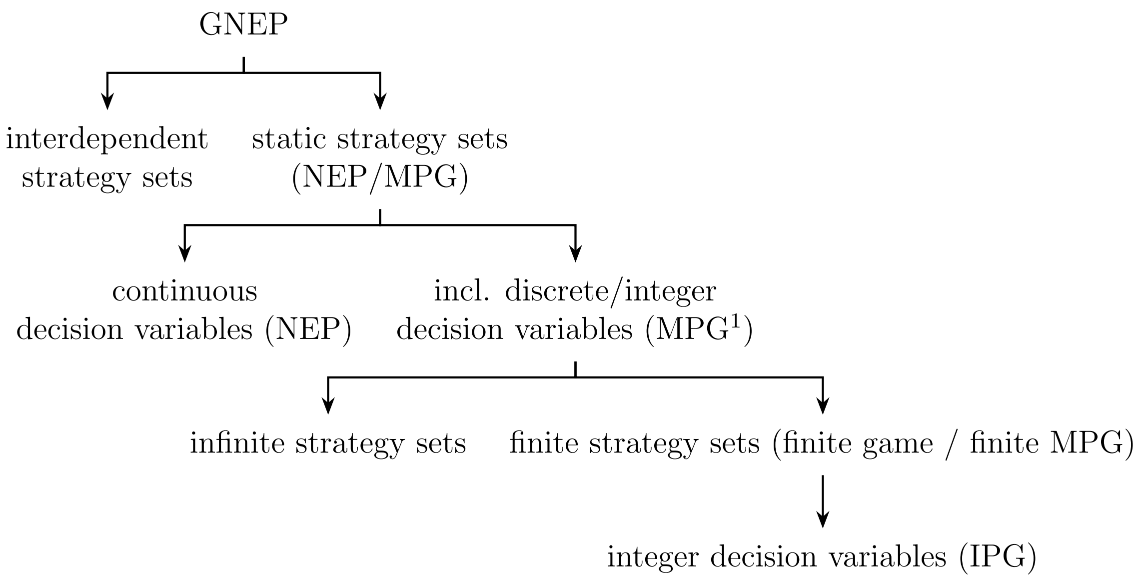

Figure 1.1: Classes of Nash equilibrium problems, based on Köppe et al. (2011) and Dragotto et al. (2021)

[1] In some (rare) cases, the term (G)NEP is also used for simultaneous games with discrete or mixed-integer decision variables.

1.2 Contribution and research questions

Building on an existing sample-generation method for the identification of a single equilibrium in IPGs (Carvalho et al., 2022), we derive a method for the identification of all equilibria and the selection of the most probable equilibrium in (complete information and non-cooperative) finite games. More specifically, we answer the following research questions (RQ):

RQ 1.1 *How can we identify all equilibria in a finite game?*

RQ 1.2 *Having identified all equilibria in a finite game – how can we select the equilibrium that is expected to be the outcome of the game based on the respective player incentives and the equilibrium selection theory by Harsanyi, 1995?*

In addition to synthetic benchmark problems such as knapsack games, competitive facility location and design problems, and competitive lot-sizing problems, we apply the derived method to a case study of hydrogen fuel station location under emerging competition. Here, we are concerned with the following research questions:

RQ 2.1 *How will competition between station providers influence the emerging hydrogen refueling network structure?*

RQ 2.2 *How valuable is it for decision makers to take competitor actions into consideration?*

RQ 2.3 *Should policymakers foster (e.g., through government-backed provider associations) or impede (e.g., through strict antitrust laws) collaboration between competing providers?*

From a practical perspective, the identification of parameters that give rise to an observed equilibrium can be equally insightful as the identification of the equilibrium itself. Third parties, such as policy makers or new entrants, might wonder why a specific equilibrium emerged, why this equilibrium can be considered stable or which parameters induced the decision makers to put in place the observed strategies. Building on an existing inverse optimization approach for non-competitive integer programs, we present a method to extract hidden information from observed equilibria in IPGs and show its application to improve decision making for a new entrant in a competitive retail location problem. Here, we address the research questions:

RQ 3.1 *How can we identify parameters that lead to observed equilibria in situations of simultaneous competition using inverse optimization, taking into account integrality constraints?*

RQ 3.2 *How valuable are the derived parameter estimates to new entrants, optimizing their own market entry location strategy?*

1.3 Complexity of equilibrium identification and selection

In the very first research question (**RQ 1.1**) we set out to identify all equilibria of a finite mathematical programming game. Yet, already finding even a single Nash equilibrium can be considered hard: Nash (1951) shows that for a finite game at least one Nash equilibrium in pure or mixed strategies exists. Building on this existence guarantee, Papadimitriou (2007) shows that the identification of a Nash equilibrium is PPAD-complete, i.e., using polynomial parity arguments on directed graphs it can be shown that the problem is guaranteed to have a solution, and a solution can be verified by a deterministic polynomial time algorithm (see Papadimitriou, 1994). However, even just determining whether a two-player game in strategic form has at least two Nash equilibria is NP-complete (Gilboa and Zemel, 1989). By definition, the identification of all equilibria in a finite game is at least as hard.

Kamal Jain has been famously quoted with the saying "if your laptop cannot find it [the equilibrium], neither can the market" (Papadimitriou, 2007), highlighting the limited predictive power of an equilibrium concept that cannot be computed by all market participants. Given the above complexity results, one might be tempted to ask why even bother devising a method addressing **RQ 1.1**. Within this thesis, we will show many practical, realistically sized problems that can be solved within time limits which are adequate for strategic decision making in practice. Yet, based on the above complexity results and preliminary experiments, it is to be expected that there are examples of competitive decision making in practice which cannot be solved within a reasonable timeframe by the methods discussed in this thesis.

Similarly, in **RQ 1.2**, we set out to select the most probable equilibrium. In Chapter 3, we reformulate this problem to several volume computation problems of non-linearly bounded sets in multiple dimensions. Note that already the volume computation of linearly bounded sets in multiple dimensions (polyhedra) is known to be #P-hard (Dyer and Frieze, 1988), where #P refers to the set of counting problems where the solution equals the number of paths to acceptance of a non-deterministic polynomial time Turing machine (Valiant, 1979). We therefore do not determine the volumes exactly and resort to numerical volume estimation based on a Monte Carlo approach.

1.4 Outline

Chapter 2 reviews related literature, addressing prior contributions on equilibrium identification and selection, inverse optimization and location selection for (alternative) fuel stations and retail outlets.

In Chapter 3, we address the equilibrium identification and selection problem for finite MPGs. We devise a column-and-row generation method, with which we sample only relevant parts of a MPG and identify equilibria within this sample. We show that, under the condition the examined MPG is finite, the set of equilibria of the final sample is equivalent to the set of all equilibria of the original game. We subsequently select the most probable equilibrium among the set of all equilibria based on an adaptation of Harsanyi's equilibrium selection theory (Harsanyi, 1995). This chapter is co-authored by Stefan Minner (Technical University of Munich) and has been made available as a working paper under Crönert and Minner (2021a).

In Chapter 4, we examine a scenario of emerging competition under hydrogen fuel station providers. We extend the Flow capturing location-allocation model (FCLM), a widely adopted flow-based location problem for alternative fuel stations, to this setting of simultaneous competition and apply concepts developed in Chapter 3 to solve the resulting model formulation. Inspired by the Value of the stochastic solution (VSS) in stochastic programming (Birge and Louveaux, 2011), we highlight the benefit of the novel approach by introducing the Value of the competitive solution (VCS). Thus, we quantify the advantage of taking into account apparent competition in comparison with a naïve approach assuming a monopoly. This chapter has been published under Crönert and Minner (2021b) and has been co-authored by Stefan Minner (Technical University of Munich).

We address the inverse optimization problem for IPGs with a novel cut-generation method in Chapter 5. We apply the derived method to a competitive retail location problem and highlight the benefit of the extracted hidden information for a new entrant observing the current location structure of incumbents. The chapter is available as a working paper under Crönert et al. (2022), co-authored by Layla Martin (Eindhoven University of Technology), Stefan Minner (Technical University of Munich) and Christopher S. Tang (University of California, Los Angeles).

Chapter 6 summarizes methodological contributions and insights. It also discusses limitations of the presented work and provides opportunities for further research.

Chapter 2

Related literature

This chapter reviews literature about related problem classes and methodological approaches (Section 2.1 and 2.2) as well as problem formulations concerning competitive location selection for alternative fuel stations and retail outlets (Section 2.3).

2.1 Equilibrium identification and selection

The structure, in which a Nash equilibrium problem is modeled, drives the choice of equilibrium identification approaches. Problems with a small number of strategies and players can be represented in normal-form, enumerating all possible strategy combinations and their associated payoffs per player in a matrix. Generalized Nash equilibrium problems (GNEPs) and Mathematical programmming game (MPG) formulations can succinctly represent larger games. Here, strategy sets and combinations are not enumerated, but represented through various (in)equality constraints. Therefore, we review contributions on equilibrium identification separately for normal-form games (2.1.1), GNEPs (2.1.2) and MPGs (2.1.3), before discussing prior work on equilibrium selection in Section 2.1.4.

2.1.1 Normal-form games

The majority of approaches able to identify Nash equilibria are targeted towards normal-form games, examples include the Lemke-Howson (Lemke and Howson, Jr., 1964) and Porter-Nudelman-Shoham (PNS, Porter et al. (2004)) algorithms as well as the MIP-Nash proposed by Sandholm et al. (2005) based on a mixed-integer programming formulation. We describe the MIP-Nash algorithm in detail in Section 3.2.2, for an overview of the Lemke-Howson algorithm and other general equilibrium computation approaches, see Stengel (2007). Algorithms requiring a normal-form representation of games (such as MIP-Nash) cannot be used to effectively address MPGs. Köppe et al. (2011) point out already the normal-form representation of such a game is exponential in size, since all possible strategy combinations between players need to be enumerated.

2.1.2 Generalized Nash equilibrium problem (GNEP)

The GNEP is a generalization of the original Nash equilibrium problem (NEP) in which the player's action set is no longer static but can instead depend on the competing players' selected strategies. For a formal definition of GNEPs and an extensive review of contributions in the field, see Facchinei and Kanzow (2007). Whereas finite MPGs and Integer programming games (IPGs) assume discrete or integer decisions, in GNEPs decision variables are usually assumed to be continuous (see, e.g. the review of Facchinei and Kanzow (2007)). Discrete or integer decisions are not adequately captured in these continuous formulations. For example, solving the continuous relaxation of an IPG does not provide any meaningful information on the number of Nash equilibria, let alone the detailed strategy set of any (mixed) equilibrium of the original IPG (Carvalho et al., 2022). However, there have been previous efforts to enable computation and selection of Nash equilibria in discrete or mixed-integer GNEPs (Hemmecke et al., 2009; Sagratella, 2017a, 2019; Sagratella et al., 2020). Note that these efforts are limited to pure strategy equilibria, which may fail to exist in general: For example, Hemmecke et al. (2009) propose an algorithm to identify pure equilibria in discrete GNEPs with monotonously decreasing payoff functions. In Sagratella (2017a) different solution methods to identify pure equilibria for mixed-integer GNEPs with a potential function are discussed. Sagratella (2019) presents an enumerative algorithm to compute the whole pure equilibrium set for GNEPs with mixed-integer variables. Subsequently, Sagratella et al. (2020) show an application of this approach to the (competitive) fixed-charge transportation problem.

2.1.3 Mathematical and integer programming games

Methods utilizing the succinct representation of large strategy sets in MPGs and IPGs are currently limited to the identification of one equilibrium out of many, or are restricted to the set of all pure equilibria. Literature on general MPGs is scarce with the term only recently being introduced by Dragotto et al. (2021). Most prior work focuses on special cases of MPGs, such as IPGs (Carvalho et al., 2022; Huppmann and Siddiqui, 2018;

Köppe et al., 2011; Sagratella, 2016, 2017b), or reciprocally-bilinear games (Carvalho et al., 2021).

Köppe et al. (2011) provide a theoretical approach to compute pure Nash equilibria for IPGs where payoff-functions are the differences between piecewise-linear convex functions. However, the authors acknowledge that the actual implementation of their algorithm and its practical application to integer programming games are still out of reach.

Sagratella (2016) proposes a branching method capable of computing the set of all pure equilibria of IPGs, which he calls Nash equilibrium problems with discrete strategy sets. Similar to its GNEP generalization (Sagratella, 2019), the approach iteratively solves the continuous relaxation of the IPG, and does not determine mixed equilibria. For a newly introduced class of "2-groups partitionable" games, the author proves existence of a (pure) equilibrium and provides an alternative, computationally advantageous Jacobi-type algorithm for the identification of this equilibrium. Sagratella (2017b) generalizes this algorithm to provide approximate (pure) equilibria for mixed-integer, "2-groups partitionable" games and presents an application for a Cournot oligopoly model.

Huppmann and Siddiqui (2018) study "binary equilibrium problems" in the context of electric power markets. In contrast to general IPGs with integer decision variables, they restrict the decision variables to be binary. Their proposed solution method is able to find pure-strategy Nash equilibria or "quasi-equilibria" in situations where no pure equilibrium exists, while neglecting any mixed equilibria. In a quasi-equilibrium no player can improve their own payoff through unilateral deviation by more than the compensation payment offered by a third party operator or regulator. Using a multi-objective program, this third party operator is assumed to select a quasi-equilibrium in which the required compensation is minimal, in combination with another solution criterion (e.g., welfare-maximization). In contrast, the approach discussed in Chapter 3 does not rely on the existence of a third party for equilibrium selection. The approach by Huppmann and Siddiqui (2018) scales well in the number of players (the authors solve instances with up to 64 players), but is limited in the number of binary decisions, as switching costs need to be determined for all possible realizations of every binary decision.

Carvalho et al. (2022) propose a Sample generation method (SGM) for IPGs. They start with a sample IPG, where players are restricted to a small subset of their original action set. As this sample is finite and sufficiently small, it can be transformed into normal-form through the enumeration of all strategy combinations in the sample, and traditional equilibrium finding approaches can be applied. The authors propose the application of PNS (Porter et al., 2004) to identify a Nash equilibrium of this sampled game. Having identified this (sample) equilibrium, they search for best-responses for each player individually. If there is a best-response strategy for any player, that gives the respective player an incentive to deviate from the identified (sample) equilibrium, they add the best-response strategy to the player's sampled strategy set and repeat the identification of an equilibrium in the new sampled game and the subsequent search for best-responses. The algorithm terminates when no best-response can be found to the Nash equilibrium of the sampled game, which proves that it is indeed an equilibrium to the original IPG. The SGM can compute a single Nash equilibrium for IPGs. While the identified equilibrium heavily depends on the initialization, the authors did not identify systematic advantages of any examined initialization procedure. With a larger initial sample (or using multiple different initializations) the SGM could be used to find additional equilibria, however, the identification of all equilibria cannot be guaranteed. In an effort to keep the sample size at a minimum, Carvalho et al. (2022) discuss an alternative to the SGM, called m-SGM. In the m-SGM, in each iteration the PNS algorithm is limited to equilibria that are supported by the newly added best-response strategies. In case such an equilibrium fails to exist, the method backtracks to the sample of the previous iteration and calculates an alternative equilibrium to that sample. Although backtracking rarely occurs in their numerical experiments, the authors report computational advantages due to the limitation of PNS to equilibria that include newly added best-response strategies.

For reciprocally-bilinear games, in which each player's objective is linear in its own variables and bilinear in its competitors' variables, Carvalho et al. (2021) propose a "cut-and-play algorithm" based on constructing increasingly tight approximations of the convex hull of each player's action set in the original game and solving the corresponding approximated game. Their approach returns a single mixed equilibrium or a proof of non-existence of mixed equilibria.

2.1.4 Equilibrium selection

The introduced algorithms for IPGs (Carvalho et al., 2022; Huppmann and Siddiqui, 2018; Köppe et al., 2011; Sagratella, 2016), or discrete GNEPs (Hemmecke et al., 2009; Sagratella, 2017b, 2019) do identify a single mixed equilibrium, or some or all pure equilibria, but cannot identify all mixed equilibria, and thereby do not fully address the equilibrium selection problem. However, some authors acknowledge the problem of multiple equilibria and extend their equilibrium identification approaches accordingly to partially address the equilibrium selection problem. For example, Sagratella (2019) proposes an approach to select an equilibrium among the set of all pure equilibria. Here, the selection criterion can be any real-valued function (defined over the equilibrium set), by consequence the selected equilibrium is not necessarily unique nor unambiguous. Similarly, Carvalho et al. (2022) discuss an extension of their SGM-approach to choose the optimal (e.g., welfare-maximizing) equilibrium among the set of all correlated equilibria.

In general, theory on equilibrium selection for non-repeated games builds on refinements of the original Nash equilibrium, ruling out any unreasonable (e.g., dominated) strategies. Following this concept, there are several approaches that enable the reduction of Nash equilibria to a smaller set of "stable" equilibria (Myerson, 1978; Selten, 1975). While the exact definition of stability differs between these approaches, they share a common drawback: They do not reduce the set of Nash equilibria to a single equilibrium, which limits their predictive power. Based on further refinements, the methods proposed by Harsanyi and Selten (1988) and Harsanyi (1995) arrive at a single selected equilibrium.

Harsanyi and Selten (1988) propose a one-point solution theory: Among a set of equilibria, they select the most stable equilibrium predominantly based on pay-off dominance and risk dominance properties. A pay-off dominant equilibrium yields higher payoffs for all players compared with a pay-off dominated alternative. Playing accordingly with a risk-dominant equilibrium yields higher expected payoffs across players compared with a risk-dominated equilibrium. This expected payoff is calculated based on the assumption that the likelihood of player p to play accordingly with an equilibrium can be measured using the ratio of p's payoff within this equilibrium compared with p's payoffs for strategies in alternative equilibria. As this stability concept does not necessarily determine a unique solution, the authors complement their selection theory with a tracing procedure that substitutes a set of (equally stable) equilibria with their centroid. Repeated elimination of dominated candidates and substitution yields a unique result in a finite number of steps.

In their original theory of equilibrium selection, Harsanyi and Selten (1988) assume that players choose the payoff-dominant equilibrium in cases where a payoff-dominant equilibrium is risk-dominated by an alternative equilibrium. As Aumann (1990) provides examples in which no rational player would choose the payoff-dominant over the risk-dominant equilibrium, Harsanyi (1995) proposes a new multilateral risk-dominance measure as the sole criterion for equilibrium selection.

Aware that there are alternative equilibrium selection theories, and that theory on equilibrium selection cannot fully cover behavioral decision making in practice (for an experimental comparison, see for example Charness et al. (2014)), we focus on Harsanyi's equilibrium selection theory when referring to the expected outcome of a game in the following.

2.2 Inverse optimization and inverse equilibrium problems

Given a (feasible) solution of an optimization problem (forward problem), the objective of inverse optimization is to find a set of parameter values such that a variation of the forward problem based on these parameters leads to optimality of the provided solution. If there are multiple parameter values that satisfy this condition, usually the parameter values requiring the smallest adjustment with regards to the original problem are selected (Heuberger, 2004).

Among other applications, inverse optimization has been employed to identify arc capacities or costs that render a feasible solution optimal in minimum cut, minimum cost flow and assignment problems (Ahuja and Orlin, 2001), to elicit supplier costs in lot-sizing problems (Egri et al., 2014; Pibernik et al., 2011), to improve the incentive design for the sharing of savings and investments among Medicare providers (Aswani et al., 2019), and to learn the route choice behavior of experienced drivers in vehicle routing problems (Chen et al., 2021).

Inverse optimization models differ in the underlying forward problem, focusing on (i) linear problems (Ahuja and Orlin, 2001), (ii) conic problems (Zhang and Xu, 2010) or (iii) (mixed-)integer programs (Bodur et al., 2022; Moghaddass and Terekhov, 2021; Schaefer, 2009; Wang, 2009). As we study the inverse optimization of IPGs, our work is most closely related to (mixed-)integer programs (iii). Here, Schaefer (2009) provides a reformulation of an inverse integer problem into a linear program, exponentially large in the number of constraints of the original inverse problem. Wang (2009) discuss a cutting-plane algorithm for inverse mixed-integer programs, which decomposes the inverse optimization problem into a (relaxed) master problem and a cut-generating subproblem. Moghaddass and Terekhov (2021) extend this to multiple, noisy observations. Additionally, Bodur et al. (2022) propose using interior points in combination with trust-regions instead of extreme points of the subproblem to speed up the generation of cuts.

Another distinction of inverse optimization models is the differentiation between (a) deterministic or (b) noisy settings (Aswani et al., 2018). Chan et al. (2021) refer to the deterministic setting as "classic inverse optimization", and call the noisy setting "data-driven inverse optimization". In a deterministic setting, only a single optimal solution is observed and used for parameter estimation. In contrast, in Chapter 5 we examine a noisy setting, allowing for multiple, approximately optimal observations. Whereas most of the aforementioned contributions (Ahuja and Orlin, 2001; Schaefer, 2009; Wang, 2009; Zhang and Xu, 2010) focus on deterministic settings, Aswani et al. (2018), Bärmann et al. (2017), and Bertsimas et al. (2015) and Moghaddass and Terekhov (2020, 2021) introduce methods that combine inverse optimization with multiple and/or noisy observations.

Lastly, inverse optimization approaches can be categorized based on their ability to capture (simultaneous) competition between multiple decision-makers. While prior applications of inverse optimization to equilibrium problems with multiple decision-makers exist, these contributions are limited to continuous decisions and cannot address IPGs. For example, Bertsimas et al. (2015) propose a data-driven estimation approach to inversely optimize parameters leading to observed Nash and Wardrop equilibria. In exemplary applications, they derive the cost functions in congestion games or demand information under Bertrand-Nash competition. Similarly, Allen et al. (2022) use inverse optimization to learn cost functions in generalized Nash games with continuous decision variables and joint constraints between players.

2.3 Location selection

Location selection is a prime example for strategic integer (or binary) decision making under competition. Location decisions are often high cost and not easily reversible, at the same time their impact on customer choice can be considerable. As such, it is crucial to include potential competitor (re-)actions into one's location selection. In the following, we differentiate between location selection for alternative fuel stations (Section 2.3.1), relying predominantly on a flow-based representation of customer demand traveling from origin to destination, and classical retail location selection (Section 2.3.2).

2.3.1 Alternative fuel station location selection

Recent literature on alternative fuel station locations focuses on the optimal location of Electric vehicle (EV) charging stations and rarely addresses the Hydrogen fuel station (HFS) location problem (Ko et al., 2017). Pagany et al. (2019) provide a comprehensive review of more than 160 contributions proposing spatial localization methodologies for the EV charging infrastructure. Guo et al. (2016) propose an equilibrium model for competitive infrastructure planning of fast charging stations. Their multi-agent optimization problem with equilibrium constraints builds on node-based demand formulations.

While the HFS location problem is related to these EV charging station applications, there are some key characteristics that differ and motivate dedicated research for HFS:

- Charging time and station capacity: Even in the best-case with DC-based fast chargers, Battery electric vehicles (BEVs) require charging times of more than 20-30 minutes whereas refueling at an HFS takes an Fuel cell electric vehicle (FCEV) only 3-5 minutes (Shoettle and Sivak, 2016). This implies that limited capacities (and queuing of vehicles) are important to consider for chargers (see e.g., Jung et al. (2014) and Upchurch et al. (2009)), but are not as crucial for hydrogen fueling stations.

- Investment costs: The investment costs for hydrogen fueling stations exceed the investment cost for charging stations by factor 10-25. Investment costs for fast DC-chargers are below 100k EUR (Schroeder and Traber, 2012), whereas investment costs for HFS range from 1 to more than 2.4mn EUR (Apostolou and Xydis, 2019). With cost implications of this magnitude, considering uncertainties in competitor actions becomes even more important for HFS location selection.

- Existing infrastructure: While charging station location selection for EVs can be limited through the accessibility of a high-voltage grid (Muratori, 2018), hydrogen fueling stations will be limited by the existence of conventional fuel stations or will be constructed jointly with conventional fuel stations as they share similar considerations with regards to safety concerns and space requirements (Dagdougui et al., 2018).

As drivers – and particularly drivers relying on alternative fuels – tend to refuel along their route (e.g., during their commute) rather than on start- or end-points of their trip (Kuby et al., 2013), fuel station location models commonly assume path-based demand. Whereas in node-based location models (e.g., p-median) demand is concentrated to distinct nodes in the network, path-based demand models assume that demand occurs along a path and any facility located along the path could serve the customer mid-route (Revelle and Eiselt, 2005). For alternative fuel stations, Upchurch and Kuby (2010) show that path-based demand models are more robust: Facilities located with a path-based demand model perform better on a node-based objective function, than facilities located with a node-based demand model do on a path-based objective function.

Facility location models accounting for path-based demand were introduced by Hodgson (1990) as an extension to the original covering location-allocation problem by Cooper (1963) and dubbed the Flow capturing location-allocation model (FCLM). The FCLM maximizes captured path-based demand by strategically placing facilities along the route of customers. Since its inception, the FCLM has been extended to account for small path deviations or detours (Berman et al., 1995a) or a limited vehicle driving range (He et al., 2018; Kuby and Lim, 2005; Kuby et al., 2009). MirHassani and Ebrazi (2013) discuss a flexible reformulation of this range-constrained FCLM (also called flow refueling location model (Kuby and Lim, 2005, FRLM)), that proves computationally advantageous for larger problem sizes. Arslan et al. (2019) develop a branch-and-cut algorithm to further improve solution times. Other extensions address stochastic customer demand (Berman et al., 1995b) and multi-period network expansions (Chung and Kwon, 2015). While some prior extensions to the FCLM (e.g., Berman and Krass (1998) and Wu and Lin (2003)) address competition, in contrast to our setting of simultaneous (emerging) competition, they assume a static competitive landscape with a priori knowledge about competitor locations.

There are several prior applications of the FCLM to the hydrogen fuel station location problem: Variations of the FCLM addressing range limitations with strategic placement of hydrogen fuel stations have been discussed (Capar and Kuby, 2012; Kim and Kuby, 2011; Kuby et al., 2009; Lim and Kuby, 2010; Upchurch and Kuby, 2010), as well as the option to account for (slight) deviations of customers from their shortest path (Huang et al., 2015; Kim and Kuby, 2011; Li et al., 2016). The majority of HFS location models approaches the problem from an infrastructure planning or policy making perspective. As such, they do not consider the competitive nature of multiple market players simultaneously entering an emerging market. Only Bersani et al. (2009) address possible competition within their framework. However, similar to the general competitive FCLM formulations (Berman and Krass, 1998; Wu and Lin, 2003), competitor locations are assumed to be known beforehand, rather than being sequentially or simultaneously decided by new entrants.

In conclusion, prior work on path-based models for hydrogen fuel station, alternative fuel station or electric charging station locations has been limited to non-competitive or simplified competitive settings where competitor locations are known a priori. However, as FCEV adoption and by consequence the competitive landscape of HFSs is still emerging, any location that is optimal today, might become sub-optimal in the short-term based on competitor actions.

2.3.2 Competitive location selection in retail applications

Competition in facility location has been studied intensively (see the reviews Drezner, 2014; Eiselt et al., 2019) since the seminal works of Hotelling (1929) and Huff (1964) who considered location choice on a line and probabilistic customer choice behavior, respectively. Commonly, competitive location models focus on optimal location decisions of one retailer and assume competitor location decisions to be static (Plastria, 2001) or to be subject to a leader-follower relationship (Eiselt and Laporte, 1997). Studies of optimal locations under simultaneous competition (see e.g., Chapter 4 or Godinho and Dias, 2010) assume full knowledge of both the competitors' profit function and the customers' store choice behavior.

To approximate this customer choice and competitive interactions from empirical data, a second stream of literature suggests to use maximum likelihood estimation (e.g., Seim, 2006; Zhu and Singh, 2009) or regression models (e.g., Shriver and Bollinger, 2022) on sales level data. For example, Seim (2006) empirically examines location decisions for video retailers in a simultaneous game of incomplete information. Zhu and Singh (2009) extend this model to allow for firm-specific competitive effects and show an application to discount retailers (Wal-Mart, Kmart, Target). They quantify strong returns for spatial differentiation, given negative impacts among competitors in close proximity. Shriver and Bollinger (2022) show that proximity to a store shifts customers from an online channel to brick-and-mortar stores and limits the chain's ability to price discriminate. However, maximum likelihood estimation and regression models require a large and detailed amount of data, such as historical sales data on the point of sales level which might not be available to outside parties (e.g., new entrants).

Similarly, parameter estimation in Huff-like models requires more data than publicly available. Leszczyc et al. (2004) estimate parameters of such a choice model using customer-specific data to explain competitive effects between retailers, differentiating between single-purpose and multi-purpose shoppers. Li and Liu (2012) use an extended Huff model to examine competition between Walmart and K-Mart based on empirical data for Cincinnati. They find that location choice drives Walmart's competitive advantage rather than an inherent brand effect. Pancras et al. (2012) calibrate a multinomial logit choice model to investigate the influence of cannibalization on facility location of fast-food restaurants in a dynamic environment. They combine sales data from a fast-food chain with demographic data to derive detailed insights into the level of cannibalization between restaurants of this chain. For the examined chain, they find that cannibalization effects between two restaurants belonging to the same chain gradually decrease with distance and are negligible at a distance of up to 10 miles. They do not account for competitive effects with stores belonging to competing chains, partly due to the lack of insight into competitors' store-level sales data.

Chapter 3

Equilibrium identification and selection in finite games

Finite games provide a framework to model simultaneous competitive decisions among a finite set of players (competitors), each choosing from a finite set of strategies. Potential applications include decisions on competitive production volumes, over capacity decisions to location selection among competitors. The predominant solution concept for finite games is the identification of a Nash equilibrium. We are interested in larger finite games, which cannot efficiently be represented in normal-form. For these games, there are algorithms capable of identifying a single equilibrium or all pure equilibria (which may fail to exist in general), however, they do not enumerate all equilibria and cannot select the most likely equilibrium.

We propose a solution method for finite games, in which we combine sampling techniques and equilibrium selection theory within one algorithm that determines all equilibria and identifies the most probable equilibrium. We use simultaneous column-and-row generation, by dividing the n-player finite game into a MIP-master problem, capable of identifying equilibria in a sample, and two subproblems tasked with sampling (i) best-responses and (ii) additional solution candidates.

We show algorithmic performance in 2- and 3-player knapsack games as well as facility location and design games, highlighting differences in solutions between the proposed approach and state of the art. Thereby, we enable decision makers in competitive scenarios to base their actions on the most probable equilibrium.

3.1 Introduction

3.1.1 Problem setting

Hardly any decision is made in isolation and in fact most decision makers are dealing with fierce competition when trying to find the optimal decision for their problem. The expected outcome of such a competitive problem setting, or the individually optimal course of action for each competitor is not evident. Finite games allow us to model the simultaneous competitive decision making between a finite set of players, where each player is restricted to a finite set of strategies (action set). We are interested in large finite games, where the size of the action set impedes the full enumeration of all possible strategy combinations among all players in normal-form. Instead, we assume the action set of each player to be represented succinctly by a set of (potentially non-linear) inequalities, payoffs are calculated based on a player specific objective function depending on the players' own decisions and decisions of their competitors. We assume a non-cooperative full-information game, where players are risk-neutral and purely self-interested, while having complete information of each other's payoffs and action sets.

As we assume finite action sets, the existence of mixed equilibria follows directly from Nash (1951). While this guarantees the existence of at least one equilibrium, it also implies the possible existence of multiple equilibria. The existence of multiple Nash equilibria is an issue of considerable practical importance (Köppe et al., 2011), leading to twofold implications: First, this makes it necessary to have a systematic methodology that can find all equilibria to a given finite game. Second, the methodology should be able to select the most probable equilibrium, as the enumeration of (all) possible equilibria holds limited practical value: A specific Nash equilibrium is a reasonable outcome of a game, only if every player knows which Nash equilibrium the competitors are intending to play (Harsanyi and Selten, 1988). In the general case of multiple Nash equilibria, this can only be the case if there is pre-play communication between players or, in non-cooperative games, if all players adhere to the same theory or principle when selecting one among multiple Nash equilibria.

The concept of finite games has many applications; among them are investment decisions in manufacturing technologies (Röller and Tombak, 1993) or location selection and supply chain design under competition (Dobson and Karmarkar, 1987; Serra et al., 1992), as well as discrete capacity games (Anderson et al., 2017) or games modeling assortment competition with multiple products (Federgruen and Hu, 2015). Within transportation and traffic science, finite games can be used to model and compute user equilibria (the steady-state traffic flow pattern, in which no traveler benefits from changing their route, see Friesz and Bernstein, 2016). In addition, finite games can be used to adequately model problems that, despite being integer by nature, are as of now predominantly discussed in their relaxed, continuous variant. Examples include decisions on (discrete) capacities, as in (simultaneous) variants of the original Spence-Dixit-Model (Dixit, 1980; Spence, 1977) or competitive newsvendor problems (Lippman and McCardle, 1997), as well as decisions on (discrete) booking class capacity levels in revenue management games among airlines (Netessine and Shumsky, 2005). Similarly, competitive inventory models (e.g., Cachon and Zipkin, 1999) usually assume continuous variables although fractions of (potentially non-divisible) goods might misrepresent reality.

3.1.2 Research question and contribution

We develop a solution approach for finite games by integrating the identification of all equilibria with the subsequent selection of the single most probable equilibrium. Thus, we answer the following two research questions (RQ):

RQ 1.1 *How can we identify all equilibria in a finite game?*

RQ 1.2 *Having identified all equilibria in a finite game - how can we select the equilibrium that is expected to be the outcome of the game based on the respective player incentives and the equilibrium selection theory by Harsanyi (1995)?*

In answering these two questions, our research adds both methodological and managerial contributions to the existing literature. Methodologically, we extend prior equilibrium identification methods (Carvalho et al., 2022; Sandholm et al., 2005) to identify not only one but all Nash equilibria of an n-player finite game (for small n). We then adapt the existing equilibrium selection theory by Harsanyi (1995), to enable its application to finite games, in which the size of the action sets impedes their full enumeration in normal-form. We highlight our contribution towards managerial decision making in competitive scenarios by illustrating for selected examples how prior approaches return unlikely equilibria.

3.2 Equilibrium identification and selection algorithm

We propose an integrated solution method, combining the identification of all Nash equilibria of a finite game in an exhaustive column-and-row generation based method with the adaption of Harsanyi's equilibrium selection theory for finite games.

3.2.1 Definition of finite games

A finite game is a decision problem between a finite number of risk-neutral, non-cooperating players $N = \{1, 2, \ldots, n\}$ that possess full information of each other's payoffs and action sets. Players decide simultaneously. In a finite game, each player $p \in N$ can choose from the non-empty and finite set of strategies X^p (action set). Let Q^p denote the number of decision variables of player p. A strategy $x^p = (x_1^p, \ldots, x_i^p, \ldots, x_{Q^p}^p) \in X^p$ of player p is a specific choice of (not necessarily integer) decision variables x_i^p, chosen from a finite set of values X_i^p: $x_i^p \in X_i^p$. Decision variables x_i^p can be real-valued, as long as this set is ensured to be finite (i.e., discrete and bounded). The finite action set X^p is defined by K^p inequalities (g_k, including the upper and lower bounds of x_i^p) in only p's decision variables x^p:

$$X^p = \{x^p \mid g_k(x^p) \leq 0,\ \forall k \in \{1, \ldots, K^p\}\} \neq \emptyset \tag{3.1}$$

This definition of X^p through a set of inequalities (3.1) enables a succinct representation of larger action sets. It also permits the usage of mathematical programming solvers in the application of the proposed methodology to applied problem settings (see Section 3.3). Different to Generalized Nash equilibrium problems (GNEPs), this action set does not depend on decisions from other players. Each player p maximizes their own payoff Π^p, a function of their own strategy x^p, and the selected strategies of their competitors x^{-p}. Here, we use $(\cdot)^{-p}$ to denote $(\cdot)$ (the respective variable) for all $\tilde{p} \in N \setminus \{p\}$. Thus, x^{-p} refers to a strategy combination across all competitors of p, and X^{-p} is the action set encompassing all competitor strategy combinations. Given $x^{-p} \in X^{-p}$, a player finds a best-response to x^{-p} by solving:

$$\max_{x^p \in X^p} \Pi^p(x^p, x^{-p}) \tag{3.2}$$

A pure profile of strategies $x = (x^p, x^{-p}) \in X = X^p \times X^{-p}$ that solves (3.2) for all players, such that no player can benefit from unilaterally deviating from x^p, is called a pure(-strategy) Nash equilibrium (Nash, 1951). Such pure equilibria may not exist for some finite games, as shown for the special case of Integer programming games (IPGs) by Carvalho et al. (2018a). We therefore introduce the following notation to account for general, mixed(-strategy) Nash equilibria: $\varphi^p = (\varphi_{x^p})_{x^p \in X^p}$ denotes a mixed strategy of player p, who chooses to play strategy x^p with probability φ_{x^p}. Equally, $\varphi_{x^{-p}}$ denotes the probability that the competitors $\tilde{p} \neq p$ each select their respective strategy $x^{\tilde{p}}$ as part of $x^{-p} = (x^{\tilde{p}})_{\tilde{p} \in N \setminus \{p\}}$, and with $\varphi^{-p} = (\varphi_{x^{-p}})_{x^{-p} \in X^{-p}}$ a resulting mixed strategy combination. The expected payoff for player p in the mixed strategy profile $\varphi = (\varphi^p, \varphi^{-p})$ is:

$$\Pi^p(\varphi) = \sum_{x^p \in X^p} \sum_{x^{-p} \in X^{-p}} \varphi_{x^p} \varphi_{x^{-p}} \Pi^p(x^p, x^{-p}) \tag{3.3}$$

Accordingly, the expected payoff for player p playing a pure strategy x^p in response to a mixed strategy combination φ^{-p} is defined as:

$$\Pi^p(x^p, \varphi^{-p}) = \sum_{x^{-p} \in X^{-p}} \varphi_{x^{-p}} \Pi^p(x^p, x^{-p}) \tag{3.4}$$

A mixed Nash equilibrium is a mixed strategy profile φ that satisfies (Nisan et al., 2007, p.55):

$$\Pi^p(\varphi) \geq \Pi^p(x^p, \varphi^{-p}), \quad \forall p \in N, \ \forall x^p \in X^p \tag{3.5}$$

3.2.2 Equilibrium identification

The proposed exhaustive Sample generation method (eSGM) builds on the Sample generation method (SGM) (of Carvalho et al., 2022): Here, an IPG is reduced to a small sample of enumerated strategies (sampled game), by restricting the strategy sets of all players to a subset of their respective action set X^p, and subsequently extended, based on identified (player-specific) best-responses to an equilibrium of this sampled game. Through the iterative identification of a sample equilibrium and best-responses to this equilibrium, the SGM converges to a Nash equilibrium of the IPG.

The m-SGM (Carvalho et al., 2022) is an extension of the SGM in which the support enumeration space is restricted to include newly added best-responses. When there is no equilibrium that includes the newly added best-response, the best-responses are discarded and an alternative equilibrium to the previous sample is computed. This approach reduces the size of the sampled game and can thus be computationally advantageous when identifying a single equilibrium. As we are interested in characterizing the full set of equilibria, and not only the equilibrium supported by the newly added best-response, this advantage would not be applicable in the eSGM. We therefore build on the standard SGM.

In contrast to the standard SGM, we require two key extensions. Jointly, these two extensions enable the identification of all Nash equilibria of a finite game. A process flow diagram of the proposed eSGM algorithm can be found in Appendix 3.C. As with the standard SGM, we start with a small sample of the original finite game. This small sample of the original finite game implies that the action of all players is restricted to the respective subset $S^p \subset X^p$ of their original strategies and the sampled strategy space $S = S^p \times S^{-p} \subset X$ is sufficiently small to allow the transformation of the sampled game into normal-form. Thereby, we initially neglect the majority of possible player strategies ($X^p \setminus S^p$). We determine the set Φ_S of all Nash equilibria of the sampled game by using an extension of the MIP-Nash algorithm (Sandholm et al., 2005) to n-player games. The extended MIP-Nash algorithm represents the master problem of the proposed approach. We refer to a sampled strategy $x^p \in S^p$ as a column of this problem, as increasing the sample by one strategy will introduce additional decision variables (representing the added strategy). As any newly added column (strategy) will also lead to additional column-dependent constraints, we refer to this approach as a combined column-and-row generation approach (Muter et al., 2013).

The master problem is a constraint satisfaction problem without objective function. The absence of an objective function means that, in contrast to classical column-generation approaches (see e.g., Barnhart et al., 1998), we are not interested in extending the master problem with columns that improve upon an incumbent solution of the master problem. Instead, in subproblem I, we find destabilizing best-responses (or prove the lack thereof) that render incumbent solutions of the master problem infeasible. Similarly, in subproblem II, we find additional candidate equilibria (or prove the lack thereof), that could represent additional solutions to the master problem.

In case any of the two subproblems are able to generate additional columns, we add the generated columns and the column-dependent rows to the master problem and re-solve the extended master problem to determine the equilibria of the increased sample game. We terminate the algorithm once no additional candidate equilibrium and no further destabilizing best-response can be identified.

The proposed master problem is a MILP for 2-player games, a MIQCP for 3-player games, and games with $n > 3$ players accordingly lead to MIP formulations with polynomial degrees of $n - 1$. For games with up to 3 players, we can solve this master problem using quadratic solvers, for practical applications with more than 3 players, non-linear solvers or reformulations would be required. This statement is true regardless of the payoff function Π^p. Rather than using the full original payoff function defined in (3.2), payoffs in the master problem are represented as input parameters.

However, the two subproblems make use of the payoff function Π^p. In subproblem I, the effect of competitor actions within Π^p is fixed. In contrast, subproblem II directly uses the payoff function of the original finite game as an objective. While we do not impose any conditions on the payoff functions Π^p in general, they clearly drive the complexity of the subproblems.

Initialization

For the SGM (and by consequence for the proposed eSGM), the choice of the initial sample S can influence the required solution times (Carvalho et al., 2022). However, in contrast to the SGM, for the eSGM the identified solution does not change with the used initialization: In Theorem 3.2, we will show that – independent from the initialization – all equilibria can be identified. Carvalho et al. (2022) propose to initialize their algorithm with strategies that are optimal in absence of competition or with strategy combinations that maximize the total welfare. In the following, we formalize these two initializations, propose an alternative initial solution, and introduce δ as a measure of competitiveness of a simultaneous (Nash) game that assesses the benefit of taking into account apparent competition.

We denote with $x_{\mathrm{LB}} = (x^p_{\mathrm{LB}}, x^{-p}_{\mathrm{LB}})$ the first initialization proposed by Carvalho et al. (2022), in which players neglect any apparent competition and choose the optimal strategy when all competitor variables are set to 0. Depending on the problem formulation this zero vector can be infeasible for the competitors. Yet, even in this case, x_{LB} can be a viable initialization, representing the (feasible) best-response of each player p to the infeasible zero vector of its competitors.

Building on this competition-neglecting initialization, we propose an alternative initialization $x_{\mathrm{UB},p}= (x^p_{\mathrm{UB},p}, x^{-p}_{\mathrm{UB},p})$, in which player p maximizes the payoff Π^p by choosing its own decision variables $x^p_{\mathrm{UB},p}$ and the competitors' decision variables $x^{-p}_{\mathrm{UB},p}$, whilst ensuring that each competitor $\tilde{p} \neq p$ receives at least the payoff of the lower bound initialization $\Pi^{\tilde{p}}(x_{\mathrm{LB}})$:

$$x_{\mathrm{UB},p} = \underset{x^p, x^{-p}}{\arg\max}\, \Pi^p(x^p, x^{-p}) \tag{3.6}$$

$$\text{s.t. } \Pi^{\tilde{p}}(x^{\tilde{p}}, x^{-\tilde{p}}) \geq \Pi^{\tilde{p}}(x_{\mathrm{LB}}) \qquad \forall \tilde{p} \neq p \tag{3.7}$$

Alternatively, Carvalho et al. (2022) propose the initialization with social welfare optimal strategies. We denote with $x_{\mathrm{WF}} = (x^p_{\mathrm{WF}}, x^{-p}_{\mathrm{WF}})$ such a welfare-maximizing strategy combination, in which players coordinate to maximize the cumulative payoff across all players:

$$x_{\mathrm{WF}} = \underset{x^p, x^{-p}}{\arg\max} \sum_{p \in N} \Pi^p(x^p, x^{-p}) \tag{3.8}$$

Commonly, the Price of Anarchy (PoA) and the Price of Stability (PoS) (see Koutsoupias and Papadimitriou, 1999; Nisan et al., 2007) are used as game-theoretic measures to determine the (in-)efficiency of equilibria or the benefit of third party (e.g., governmental) involvement in non-cooperative games. As we are interested in the necessity of a game theoretic formulation (i.e., how important it is for competitors to recognize competition), we propose to extend the PoA and the PoS with a measure of competitiveness based on the introduced initializations ($x_{\mathrm{LB}}, x_{\mathrm{UB},p}$):

Definition 3.1 (Competitiveness of a non-cooperative game)**.**
We define the competitiveness δ of a non-cooperative, full information game, as the average difference in payoffs across all players between the initializations ($x_{\mathrm{LB}}, x_{\mathrm{UB},p}$)*:*

$$\delta = \frac{1}{n} \sum_{p \in N} \frac{\Pi^p(x_{\mathrm{UB},p}) - \Pi^p(x_{\mathrm{LB}})}{\Pi^p(x_{\mathrm{UB},p})} \tag{3.9}$$

$$0 \leq \delta \leq 1 \tag{3.10}$$

Without loss of generality we assume positive payoffs. The more beneficial it is for players to react to apparent competition, the larger the enumerator and the closer δ to 1. In a non-competitive game, where no player can gain an advantage by deciding on behalf of competitors, δ equals 0. In contrast, in a fully competitive game, where neglecting competition would yield no payoff at all, δ equals 1.

Computation of all Nash equilibria for small finite games (master problem)

The sufficiently small size of the sampled game enables the use of general (normal-form based) algorithms for equilibrium identification. In contrast to the PNS algorithm (Porter et al., 2004) adopted by Carvalho et al. (2022), we opt for the application of MIP-Nash. While PNS shows comparable performance when only a single equilibrium is computed, MIP-Nash has computational advantages when identifying all equilibria of a (sampled) game (Sandholm et al., 2005).

The MIP-Nash algorithm proposed by Sandholm et al. (2005) determines all Nash equilibria of a normal-form game as the set of feasible solutions to a combination of constraints that ensures zero regret (i.e. no player can benefit from unilateral changes) for all played strategies. The authors discuss options to further guide the search for an equilibrium, for example by adding an objective function that maximizes the total welfare (to find the social-welfare maximizing equilibrium). As we are interested in computing all equilibria of a sample, we use the constraint formulation without additional objective function. While the original formulation is given for two player games, we extend the approach to n-player games in the following.

Recall that φ_{x^p} refers to the probability of player p playing strategy x^p in a mixed strategy profile φ and $\Pi^p(x^p, x^{-p})$ denotes player p's payoff in a pure strategy profile (x^p, x^{-p}). We introduce π_{x^p} as the expected payoff for player p playing the pure strategy x^p taking into account the respective competitor reactions and their probabilities, $\overline{\pi}^p = \max_{x^p \in S^p}(\pi_{x^p})$ as the maximum expected payoff across all x^p in the sample S^p and e_{x^p} as the regret for playing x^p. Lastly, b_{x^p} is an auxiliary binary variable, set to zero for all strategies x^p with a positive probability ($\varphi_{x^p} > 0$) and set to one for all strategies that are not played ($\varphi_{x^p} = 0$).

$$\sum_{x^p \in S^p} \varphi_{x^p} = 1 \qquad \forall p \in N \tag{3.11}$$

$$\pi_{x^p} = \sum_{x^{-p} \in S^{-p}} \left(\Pi^p(x^p, x^{-p}) \prod_{\tilde{p} \neq p} \varphi_{x^{\tilde{p}}} \right) \qquad \forall p \in N,\ \forall x^p \in S^p \tag{3.12}$$

$$\overline{\pi}^p \geq \pi_{x^p} \qquad \forall p \in N,\ \forall x^p \in S^p \tag{3.13}$$

$$e_{x^p} = \overline{\pi}^p - \pi_{x^p} \qquad \forall p \in N,\ \forall x^p \in S^p \tag{3.14}$$

$$\varphi_{x^p} \leq 1 - b_{x^p} \qquad \forall p \in N,\ \forall x^p \in S^p \tag{3.15}$$

$$e_{x^p} - M b_{x^p} \leq 0 \qquad \forall p \in N,\ \forall x^p \in S^p \tag{3.16}$$

$$\varphi_{x^p} \in [0,1], b_{x^p} \in \{0,1\},\ \pi_{x^p}, \overline{\pi}^p, e_{x^p} \in \mathbb{R}_{\geq 0} \qquad \forall p \in N,\ \forall x^p \in S^p \tag{3.17}$$

Constraint (3.11) ensures that the sum of the probabilities φ_{x^p} across all strategies x^p of player p amounts to 1. The main difference to the original 2-player formulation of MIP-Nash by Sandholm et al. (2005) lies in (3.12). Here, the expected payoff π_{x^p} of a single strategy x^p is determined as the expected value over all possible competitor reactions and their probabilities. A competitor strategy x^{-p} represents a combination of individual decisions by all $n-1$ competitors. Therefore, the probability $\varphi_{x^{-p}}$ is determined as the combined probability of all competitors $\tilde{p} \neq p$ choosing their respective strategy $x^{\tilde{p}}$. We can calculate this combined probability as the product of the individual probabilities $\varphi_{x^{\tilde{p}}}$. Constraint (3.13) defines $\overline{\pi}^p$ as the maximum payoff for payer p. We calculate the regret $e_{x^p} = \overline{\pi}^p - \pi_{x^p}$ as the difference between this maximum payoff and the payoff for strategy x^p in (3.14). Constraint (3.15) defines the binary auxiliary variable b_{x^p}, which is equal to 0 for played and 1 for unplayed strategies. Indicator constraint (3.16) combines this auxiliary variable with the regret e_{x^p} to ensure that only strategies with zero regret are played (i.e., have a positive probability φ_{x^p}).

If one player could benefit from unilaterally changing their own strategy, the regret for this strategy would be larger than zero. By definition, the restriction to strategies with zero regret leads to a Nash equilibrium, where no player can benefit from a unilateral change. Hence, the set of solutions to Constraints (3.11)-(3.16) constitutes the set of Nash equilibria for an n-player game. For sample S, we denote this set of equilibria of the sampled game by Φ_S. Note that MIP-Nash requires games in normal-form, where the payoff $\Pi^p(x^p, x^{-p})$ is explicitly given for every strategy combination. As already the calculation of Π^p for all strategy combinations of a large finite game would be intractable, MIP-Nash cannot be directly applied to (unsampled) large finite games.

Identification of best-response solutions (subproblem I)

While any $\varphi \in \Phi_S$ identified by the master problem (3.11)-(3.16) is guaranteed to be a Nash equilibrium for the sampled game, this only holds true for the original finite game if we can prove the non-existence of not yet sampled best-responses for all players in the solution space of the original finite game.

To identify these best-responses (or prove the lack thereof) we solve the finite game independently for each player p, while keeping strategies for all other players fixed to $\varphi = (\varphi^p, \varphi^{-p})$. We adapt (3.2) accordingly and solve the objective function:

$$\max_{x^p \in X^p \setminus S^p} \Pi^p(x^p, \varphi^{-p}) \tag{3.18}$$

under the condition that the identified best-response has to yield a payoff higher or equal to the current (sample) equilibrium $\varphi \in \Phi_S$:

$$\Pi^p(x^p, \varphi^{-p}) \geq \Pi^p(\varphi^p, \varphi^{-p}) \tag{3.19}$$

Player p's solution space X^p remains unchanged to the original finite game (3.1). We solve (3.18) for all players $p \in N$ and for all $\varphi \in \Phi_S$ of the sampled game and collect solutions (best-responses) in the set X^p_{b}. These best-responses extend the master problem by columns and column-dependent rows.

Identification of candidate equilibria (subproblem II)

We further extend the sampled game S based on candidate equilibria that are not yet part of the sampled game but could represent equilibria of the original finite game. X_c denotes the set of all candidate equilibria, with $x_\mathrm{c} = (x_\mathrm{c}^p, x_\mathrm{c}^{-p}) \in X_\mathrm{c}$ as a candidate equilibrium. To keep this set and by consequence the sample size small, we need to find conditions for candidate equilibria that are fulfilled by all Nash equilibria, while limiting the number of added non-equilibria.

From the definition of a Nash equilibrium, it becomes evident that every pure strategy that is part of the support of an equilibrium has to be optimal for the respective player with regards to unilateral changes. This observation underlies both the MIP-Nash algorithm (Sandholm et al., 2005) and the PNS algorithm (Porter et al., 2004) and is formalized in (3.20), where x_c^p is a pure strategy in player p's support of the equilibrium φ (i.e., $\varphi_{x_\mathrm{c}^p} > 0$):

$$\Pi^p(x_\mathrm{c}^p, \varphi^{-p}) \geq \Pi^p(x^p, \varphi^{-p}) \quad \forall x_\mathrm{c}^p \in \{\tilde{x}^p \in X^p \mid \varphi_{\tilde{x}^p} > 0\},\ \forall x^p \in X^p,\ \forall p \in N \tag{3.20}$$

For a formal proof and introduction of this best-response condition, see Proposition 3.1 in Nisan et al. (2007, p.55). As testing condition (3.20) for all possible $x^p \in X^p$ would essentially imply the complete enumeration of the finite game, we propose (3.21) and (3.22) as (necessary, but not sufficient) proxy conditions: Any solution $x_\mathrm{c} = (x_\mathrm{c}^p, x_\mathrm{c}^{-p})$ to the original finite game, where any player p would benefit from unilaterally deviating from an admissible strategy, that is not yet sampled $x_\mathrm{c}^p \in X^p \setminus S^p$, to one of the already sampled strategies $x^p \in S^p$, cannot be part of a Nash equilibrium and will hence not be considered as a candidate equilibrium.

$$\Pi^p(x_\mathrm{c}^p, x_\mathrm{c}^{-p}) \geq \Pi^p(x^p, x_\mathrm{c}^{-p}) \qquad \forall x^p \in S^p,\ \forall p \in N \tag{3.21}$$

$$x_\mathrm{c}^p \in X^p \setminus S^p,\ x_\mathrm{c}^{-p} \in X^{-p} \setminus S^{-p} \tag{3.22}$$

We use (3.21) and (3.22) in combination with the solution space of the original finite game (3.1). Note that for a given player $\hat{p}$ and a combination of competitor strategies $x_\mathrm{c}^{-\hat{p}}$ we are only interested in the best-response of $\hat{p}$ satisfying (3.21). We therefore solve the problem:

$$\max \Pi^{\hat{p}}(x_\mathrm{c}^{\hat{p}}, x_\mathrm{c}^{-\hat{p}}) \tag{3.23}$$

Once a solution $\hat{x}_c$ is found, we add that solution to the set of candidate equilibria X_c. We determine all candidate equilibria by eliminating every identified solution ($\hat{x}_c = (\hat{x}_c^{\hat{p}}, \hat{x}_c^{-\hat{p}})$) from the solution space through the use of cuts (3.24) until no further solution can be found:

$$x_c^{-\hat{p}} \neq \hat{x}_c^{-\hat{p}} \qquad \forall x_c^{-\hat{p}} \in X^{-\hat{p}} \setminus S^{-\hat{p}} \tag{3.24}$$

Here, the cuts (3.24) not only restrict new solutions to be different from $\hat{x}_c$ itself, but remove all possible combinations of $x_c^{\hat{p}}$ with the competitor actions $\hat{x}_c^{-\hat{p}}$ from the solution space. This is possible, as any alternative reaction $x_c^{\tilde{p}}$ to $\hat{x}_c^{-\hat{p}}$ that would lead to a Nash equilibrium would be sampled during the identification of best-responses. We use additional cuts to reduce the solution space with each identified candidate solution $\hat{x}_c$ even further: In a new candidate equilibrium, no player can benefit from unilateral deviation to an already sampled strategy (3.21), just as no player can benefit from unilateral deviation to a previously identified candidate solution $\hat{x}_c$ that is not yet in the sample S. For each identified solution $\hat{x}_c$, we add:

$$\Pi^p(x_c^p, x_c^{-p}) \geq \Pi^p(\hat{x}_c^p, x_c^{-p}) \qquad \forall p \in N \tag{3.25}$$

Jointly, (3.24-3.25) can significantly reduce the solution space – and therefore the time for the identification of all candidate equilibria X_c – with each identified candidate solution $\hat{x}_c$.

Problem (3.23) guarantees optimality (with regards to $\Pi^{\hat{p}}$) of all candidate equilibria for one player $\hat{p}$. For all other players $p \neq \hat{p}$, X_c might include multiple solutions x_c^p for the same competitor strategies x_c^{-p}. In these cases, we subsequently reduce the set X_c to include only those x_c^p that yield higher payoffs for the respective player p in response to x_c^{-p}.

While the described search for new candidates is formulated for pure strategies, it fully covers mixed strategy Nash equilibria, as Theorem 3.1 shows. This guarantees that all pure strategies that are in the support of a pure or mixed equilibrium will be sampled and by consequence the application of MIP-Nash will yield all equilibria (see Theorem 3.2).

Theorem 3.1. *Any pure strategy x_c^p in the support of a mixed strategy Nash equilibrium, that is not already part of the sample S, satisfies constraint (3.21) and will thus be part of X_c.*

Proof. Suppose there is a pure strategy profile (x_c^p, x_c^{-p}) that does not satisfy (3.21) and is not yet part of the sample ($x_c^p \notin S^p$, $x_c^{-p} \notin S^{-p}$), but is in the support of a mixed strategy Nash equilibrium $\varphi = (\varphi^p, \varphi^{-p})$ (i.e., x_c^p is played with a positive probability $\varphi_{x_c^p}$ in φ^p by player p and x_c^{-p} is played with a positive probability $\varphi_{x_c^{-p}}$ in φ^{-p} by p's competitors). Based on the assumed violation of (3.21), there would have to be a strategy $\tilde{x}^p$ with $\Pi^p(\tilde{x}^p, x_c^{-p}) > \Pi^p(x_c^p, x_c^{-p})$ that would yield a higher payoff for player p. By consequence, we can rule out favorable payoffs for p if competitors choose to play x_c^{-p} as a potential reason for x_c^p being in the support of φ. I.e., there has to be another rationale for p to play x_c^p. Thus, x_c^p has to be a best-response to an alternative strategy $\hat{x}_c^{-p}$ played with positive probability $\varphi_{\hat{x}_c^{-p}}$ in the same equilibrium. At the same time $\hat{x}_c^{-p}$ cannot be part of the sample S^{-p}, because $\hat{x}_c^{-p} \in S^{-p}$ would imply $x_c^p \in S^p$ based on the prior sampling of best-responses. As x_c^p is a best-response to $\hat{x}_c^{-p}$, and $x_c^p \notin S^p, \hat{x}^{-p} \notin S^{-p}$, taken together $(x_c^p, \hat{x}_c^{-p})$ satisfy (3.21) and x_c^p will be sampled in X_c. □

The proof of Theorem 3.1 makes no assumption on the existence of pure equilibria. Hence, it also guarantees that all pure strategies in the support of a mixed equilibrium will be sampled, even if there is no pure equilibrium (and by consequence the corresponding mixed equilibrium will be identified subsequently using MIP-Nash).

Termination and optimality

There are two termination criteria that have to be fulfilled at the same time for the algorithm to terminate:

$$X_b^p = \emptyset \qquad \forall p \in N \tag{3.26}$$

$$X_c = \emptyset \tag{3.27}$$

The algorithm terminates when there is no best-response to any of the identified equilibria that is not yet sampled (3.26), and when no additional candidate equilibrium that is not already part of the sampled game S can be found (3.27). To verify the non-existence of additional best-responses or candidate equilibria, we show infeasibility of the column-generating subproblems (3.18-3.19) and (3.21-3.24) respectively. When both termination criteria are met, the algorithm yields all equilibria Φ of the original finite game ($\Phi = \Phi_S$).

Theorem 3.2. *Upon termination, the eSGM algorithm yields the full set of Nash equilibria (i.e., all pure equilibria and all mixed Nash equilibria with unique support in pure strategies) to the original finite game.*

Proof. Suppose that an identified solution $\varphi \in \Phi$ is not a Nash equilibrium of the original finite game. By definition, at least one player would benefit from unilaterally deviating. For this player, there would have to be a best-response that ensures a higher pay-off and the first termination criterion (3.26) would be violated. Further suppose that there is an equilibrium that is not part of the solution set Φ, with a pure strategy support that is different from all other identified equilibria in Φ. By consequence, at least one strategy x^p in the support of this equilibrium would not be part of the sampled game S. As this strategy x^p would satisfy (3.21) and (3.22), the set X_c^p would not be empty and the second termination criterion (3.27) is violated. □

As we defined X to be finite, the algorithm terminates after a finite number of steps. Note that this convergence is guaranteed in cases where conventional best-response dynamics fail to converge (e.g., games with conflicting congestion effects, see Feldman and Tamir, 2012). While the repeated (iterative) improvement steps of players might cycle in the case of best-response dynamics, the proposed eSGM (as well as the original SGM) determines best-responses simultaneously for all players and terminates in case they are already part of the prior sample, thus preventing cycles.

3.2.3 Equilibrium selection

In the following, we build on the equilibrium selection theory of Harsanyi (1995), ensuring its two main principles: The probability of a player p to choose a certain strategy is higher compared with the probability of an alternative strategy, if the former is a best-response to a larger number of potential (mixed) strategy combinations of p's competitors than the latter (i). For non-degenerate unanimity games (i.e., games in which players receive a payoff only if they successfully coordinate a common strategy), the probability of a player choosing a strategy is proportional to the strategies' payoff (ii). We operationalize these two principles for finite games where action sets are defined through inequalities, avoiding the evaluation of the first principle (i) for all possible (mixed) strategy combinations.

After termination, the eSGM yields all equilibria of the finite game and the final sample S, which includes all pure strategies in the support of any equilibrium. Since no player would choose a pure strategy that is not in the support of any equilibrium, we can reduce the sample to include only non-dominated strategies supporting an equilibrium.

We refer to this as the reduced sampled game. For each strategy in the reduced sampled game, we construct stability sets. A stability set of a strategy is the set of (mixed) strategy combinations across competitors to which the respective strategy is a best-response (Harsanyi, 1995). We formalize Harsanyi's definition of stability sets in the following. After transforming the stability sets into sets of risk indicators, in accordance with principles (i) and (ii) outlined above, the probability that a player selects a pure strategy is proportional to the size of these transformed stability sets. For each mixed equilibrium, we combine these pure strategy probabilities for all strategies in the support of the mixed equilibrium across all players to determine the equilibrium probability and thereby the most probable equilibrium.

Reduced sampled game

Harsanyi (1995) applies his equilibrium selection theory to a reduced game, obtained from the original game after the removal of inferior and duplicate strategies. We operationalize this reduction for (large) finite games, starting with the set of all equilibria and thus avoiding the transformation of the finite game into normal-form and the implied complete enumeration of all strategy combinations.

Based on the introduced eSGM-algorithm, we have all pure strategies in the support of any equilibrium of the original finite game. We denote this set $\tilde{S}_{\mathrm{r}}^{p}$, restricting the final sample S^p to pure strategies x^p with a strictly positive probability φ_{x^p} in some equilibrium $\varphi \in \Phi$:

$$\tilde{S}_{\mathrm{r}}^{p} = \left\{ x^p \in S^p \;\middle|\; \exists \left(\varphi \in \Phi,\ \varphi_{x^p} > 0 \right) \right\} \subseteq S^p \tag{3.28}$$

As no player has an incentive to unilaterally deviate from any strategy in this set $\tilde{S}_{\mathrm{r}}^{p}$, we restrict the selection approach to strategies in $\tilde{S}_{\mathrm{r}}^{p}$. We further reduce $\tilde{S}_{\mathrm{r}}^{p}$ to S_{r}^{p} through the elimination of dominated strategies within this set, leading to the reduced sampled game $S_{\mathrm{r}} = S_{\mathrm{r}}^{p} \times S_{\mathrm{r}}^{-p}$:

$$S_{\mathrm{r}}^{p} = \left\{ x^p \in \tilde{S}_{\mathrm{r}}^{p} \;\middle|\; \exists \left(x^{-p} \in \tilde{S}_{\mathrm{r}}^{-p} \mid \Pi^p(x^p, x^{-p}) \geq \Pi^p(\tilde{x}^p, x^{-p}) \;\; \forall \tilde{x}^p \in \tilde{S}_{\mathrm{r}}^{p} \right) \right\} \subseteq \tilde{S}_{\mathrm{r}}^{p} \tag{3.29}$$

Here, (3.29) ensures that for each strategy $x^p \in S_{\mathrm{r}}^{p}$ and for any alternative strategy $\tilde{x}^p \in \tilde{S}_{\mathrm{r}}^{p}$, there exists at least one competitor strategy combination $x^{-p} \in \tilde{S}_{\mathrm{r}}^{-p}$ for which $\tilde{x}^p$ does not outperform x^p (i.e., x^p is not dominated by $\tilde{x}^p$).

Probability vectors and stability sets

The probability vector $q^p = (q^p_{x^{-p}})_{x^{-p} \in S_\mathrm{r}^{-p}}$ consists of $|S_\mathrm{r}^{-p}|$ probabilities $q^p_{x^{-p}}$, which a player p assigns to its competitors choosing any pure strategy combination $x^{-p} \in S_\mathrm{r}^{-p}$ (Harsanyi, 1995). For a given strategy x^p of player p, the stability set V_{x^p} is defined as the set of probability vectors q^p, to which strategy x^p of player p is a best-response. V_{x^p} is a polytope within a simplex in $|S_\mathrm{r}^{-p}| - 1$ dimensions:

$$V_{x^p} = \left\{ q^p = (q^p_{x^{-p}})_{x^{-p} \in S_\mathrm{r}^{-p}} \middle| \sum_{x^{-p} \in S_\mathrm{r}^{-p}} q^p_{x^{-p}} \Pi^p(x^p, x^{-p}) \geq \sum_{x^{-p} \in S_\mathrm{r}^{-p}} q^p_{x^{-p}} \Pi^p(\tilde{x}^p, x^{-p}) \; \forall \tilde{x}^p \in S_\mathrm{r}^p \right\} \tag{3.30}$$

$$q^p_{x^{-p}} \geq 0 \qquad \forall x^{-p} \in S_\mathrm{r}^{-p} \tag{3.31}$$

$$\sum_{x^{-p} \in S_\mathrm{r}^{-p}} q^p_{x^{-p}} = 1 \tag{3.32}$$

Here, (3.30) formalizes the best-response condition for probability vectors q^p in the stability set V_{x^p}, (3.31) limits the probabilities of the vector q^p to non-negative values and (3.32) ensures that all probability vectors are normalized to sum to 1. This notation for V_{x^p} through a combination of linear inequalities is commonly referred to as the boundary representation of a polytope (Henk et al., 2017).

Risk indicators

Intuitively, the size of the stability set V_{x^p} in the reduced sampled game should be proportional to player p's probability to play strategy x^p. The bigger the stability set of x^p, the more likely x^p is a best-response to a given strategy of competitors. In accordance with Harsanyi (1995), we do not directly measure the size of this stability set, as such a measure would not fulfill the proportionality requirement for non-degenerate unanimity games. To satisfy this requirement, Harsanyi (1995) introduces the strictly positive vector $\tilde{q}^p = (\tilde{q}^p_{x^{-p}})_{x^{-p} \in S_\mathrm{r}^{-p}}$ and the risk indicator $r^p = (r^p_{x^{-p}})_{x^{-p} \in S_\mathrm{r}^{-p}}$ with:

$$\tilde{q}^p_{x^{-p}} = \begin{cases} q^p_{x^{-p}} & \text{, if } q^p_{x^{-p}} > 0 \\ \epsilon & \text{, else} \end{cases} \tag{3.33}$$

$$r^p_{x^{-p}} = \frac{\frac{1}{\tilde{q}^p_{x^{-p}}}}{\sum\limits_{\hat{x}^{-p} \in S_{\mathrm{r}}^{-p}} \frac{1}{\tilde{q}^p_{\hat{x}^{-p}}}} \tag{3.34}$$

With ϵ being a sufficiently small number, this indicator formally defines the function $\mu^p : \tilde{q}^p \mapsto r^p$. Note that a high value of $r^p_{x^{-p}}$ for a certain strategy combination x^{-p} indicates that the probability $\tilde{q}^p_{x^{-p}}$ is small. In other words, selecting a best-response to x^{-p} is risky for p in the sense that it requires a high payoff to offset the low probability of x^{-p} for being selected by competitors.

Incentive measurement

We apply μ^p to all probability vectors q^p in the stability set V_{x^p} after transforming the probability vectors to their strictly positive equivalent (3.33) to arrive at the image set W_{x^p}:

$$W_{x^p} = \{\mu^p(\tilde{q}^p) | \, \forall q^p \in V_{x^p}\} \tag{3.35}$$

Based on Harsanyi (1995), a strategy x^p with a larger set W_{x^p} is less risky (i.e., it is a best-response to more competitor strategy combinations) compared with an alternative strategy $\tilde{x}^p$ with a smaller set $W_{\tilde{x}^p} : \lambda(W_{\tilde{x}^p}) < \lambda(W_{x^p})$, where $\lambda(\cdot)$ denotes the volume of a set in its multi-dimensional space (Lebesgue measure). Building on this observation, Harsanyi (1995) refers to the size λ of this set W_{x^p} as the incentive of player p to play x^p. This definition fulfills principles (i) and (ii) of Harsanyi's equilibrium selection theory.

Note that W_{x^p} is a non-linearly bounded set in $|S_{\mathrm{r}}^{-p}| - 1$ dimensions. We determine the size of this set $\lambda(W_{x^p})$ – representing the incentive of p to play x^p – using a Monte Carlo approach (for details, see Appendix 3.A).

The incentive $\Psi^p_{\varphi^p}$ of p to choose the mixed strategy φ^p is the average of $\lambda(W_{x^p})$ across all pure strategies in the support of φ^p, weighted by their respective probabilities $\varphi^p_{x^p}$:

$$\Psi^p_{\varphi^p} = \sum_{x^p \in S_{\mathrm{r}}^p} \varphi^p_{x^p} \lambda(W_{x^p}) \tag{3.36}$$

The probability of a player p selecting φ^p as its selected strategy profile Θ^p is proportional to the incentive $\Psi^p_{\varphi^p}$ for said profile (Harsanyi, 1995):

$$\mathrm{Prob}(\Theta^p = \varphi^p) = \frac{\Psi^p_{\varphi^p}}{\sum_{(\hat{\varphi}^p, \hat{\varphi}^{-p}) \in \Phi} \Psi^p_{\hat{\varphi}^p}} \tag{3.37}$$

As we examine non-cooperative settings, p's probability $\text{Prob}(\Theta^p = \varphi^p)$ for φ^p is independent from competitor $\tilde{p}$'s probability $\text{Prob}(\Theta^{\tilde{p}} = \varphi^{\tilde{p}})$ for $\varphi^{\tilde{p}}$. Therefore, the probability of a given Nash equilibrium $\varphi = (\varphi^p, \varphi^{-p})$ to be the most likely outcome of the game $(\Theta = (\Theta^p, \Theta^{-p}))$ is the product of the probability of p and p's competitors choosing the respective mixed strategies that define the equilibrium φ:

$$\text{Prob}(\Theta = \varphi) = \prod_{p \in N} \text{Prob}(\Theta^p = \varphi^p) \tag{3.38}$$

The solution to a finite game is the equilibrium ξ with the highest probability:

$$\xi = \arg\max_{\varphi \in \Phi} \left(\text{Prob}(\Theta = \varphi)\right) \tag{3.39}$$

In cases of a non-unique solution to (3.39), Harsanyi (1995) suggests to use the equal probability mixture of all equilibria fulfilling (3.39) when there is pre-play communication between players, and refers to the tracing procedure (Harsanyi and Selten, 1988) in problems without pre-play communication.

3.2.4 Pareto-optimal and welfare-maximizing equilibria

Based on the general eSGM-algorithm introduced in Section 3.2.2, we propose two extensions for particular problem settings. First, in cases where some form of communication or coordination between players is possible, players might focus solely on equilibria that are Pareto-optimal. Second, in some applications there might be a third party or a transaction mechanism that ensures players are restricted to welfare-maximizing equilibria.

Pareto-optimal equilibria (eSGM-PO)

We adapt the search for additional candidate equilibria described in Section 3.2.2. In addition to the constraints (3.21)-(3.22), candidate equilibria (x_c^p, x_c^{-p}) need to fulfill:

$$\Pi^p(x_c^p, x_c^{-p}) + z_i^p L^p \geq \Pi^p(\varphi_i^p, \varphi_i^{-p}) \qquad \forall (\varphi_i^p, \varphi_i^{-p}) \in \Phi_S,\ \forall p \in N \tag{3.40}$$

$$\sum_{p \in N} z_i^p \leq n - 1 \qquad \forall i \in \{1, \dots, |\Phi_S|\} \tag{3.41}$$

$$z_i^p \in \{0, 1\} \qquad \forall i \in \{1, \dots, |\Phi_S|\},\ \forall p \in N \tag{3.42}$$

In (3.40), we set the auxiliary binary variable z_i^p whenever the payoff for player p based on the candidate strategy combination (x_c^p, x_c^{-p}) is inferior to the payoff the player would receive in an alternative equilibrium i. Here, L^p is a sufficiently large constant; without loss of generality, we can assume positive payoffs ($\Pi^p \geq 0$) and use $L^p = \max_{\varphi_i \in \Phi_S}(\Pi^p(\varphi_i))$ to provide a tight bound. Using (3.41), we ensure that in every Pareto-optimal equilibrium candidate, at least one of the n players is better off than with all other existing alternative equilibria. Constraint (3.42) restricts z_i^p to binary values.

As soon as an equilibrium is identified within the sample using MIP-Nash, we add additional cuts to ensure that the final solution set does not include Pareto-dominated equilibria: Let $(\tilde{\Pi}^p, \tilde{\Pi}^{-p})$ denote the players' payoffs in an intermediate solution to the MIP-Nash constraints (3.11)-(3.16). Any Pareto-optimal solution may not be dominated by this intermediate solution and thus we can add the following constraints:

$$\overline{\pi}^p + z^p L^p \geq \tilde{\Pi}^p \qquad \forall p \in N \tag{3.43}$$

$$\sum_{p \in N} z^p \leq n - 1 \tag{3.44}$$

$$z^p \in \{0, 1\} \qquad \forall p \in N \tag{3.45}$$

Welfare-maximizing equilibria (eSGM-WM)

If we are only interested in additional candidate equilibria with a higher social welfare, we can extend constraints (3.21)-(3.22) by:

$$\sum_{p \in N} \Pi^p(x_c^p, x_c^{-p}) \geq \sum_{p \in N} \Pi^p(\varphi^p, \varphi^{-p}) \qquad \forall (\varphi^p, \varphi^{-p}) \in \Phi_S \tag{3.46}$$

Additionally, as proposed by Sandholm et al. (2005), we limit the search of equilibria within the sampled game towards welfare-maximizing equilibria. We adapt the n-player constraint formulation of MIP-Nash by supplementing constraints (3.11)-(3.16) with the objective function:

$$\max_{\varphi_{x^p}} \sum_{p \in N} \overline{\pi}^p \tag{3.47}$$

3.3 Applications and managerial insights

We implement the eSGM algorithm and its variants in Python 3 in combination with Gurobi 9.5. The CPU times are based on an eight-core processor with 2.6 GHz base frequency. We initialize the eSGM with the player's optimal strategies when they are alone in the game (x_{LB}).

We show application of the eSGM in a knapsack game (Carvalho et al., 2022, Section 3.3.1, based on) and a large-scale facility location and design problem (Section 3.3.2). An additional application for a purely discrete version of the uncapacitated lot-sizing game proposed by Carvalho et al. (2022) can be found in Appendix 3.D.

3.3.1 Knapsack game

Problem description

Suppose that n players have to choose between a set of I technologies. Players can realize profits v_i^p for a given technology $i \in I$ and experience synergies ($c_i^{p\tilde{p}} > 0$) or penalties ($c_i^{p\tilde{p}} < 0$) when a competitor $\tilde{p} \neq p$ invests in technology i. x_i^p denotes the (binary) decision of player p to invest into technology i. Each investor maximizes individual profit:

$$\max_{x_i^p} \sum_{i \in I} v_i^p x_i^p + \sum_{\tilde{p} \neq p} \sum_{i \in I} c_i^{p\tilde{p}} x_i^p x_i^{\tilde{p}} \tag{3.48}$$

The investment costs for player p and technology i is w_i^p, the players are limited by their investment budget W^p:

$$\sum_{i \in I} x_i^p w_i^p \leq W^p \qquad \forall p \in N \tag{3.49}$$

$$x_i^p \in \{0, 1\} \qquad \forall p \in N,\ \forall i \in I \tag{3.50}$$

Computational results

We explore problem setups with $n = 2$ and $n = 3$ players, and up to $|I| = 60$ technologies. For each setup, we randomly generate 10 instances, where $c_i^{p\tilde{p}}$ and v_i^p are independently drawn from a uniform distribution on interval $[-100, 100] \cap \mathbb{Z}$. We draw w_i^p from the uniform distribution on interval $[0, 100] \cap \mathbb{Z}$ and set the budgets W^p based on the respective instance number $m \in \{0, \dots, 9\}$:

$$W^p = \left\lfloor \frac{m}{11} \sum_{i \in I} w_i^p \right\rfloor \tag{3.51}$$

Figure 3.1 summarizes the results of all game instances for $n = 2$ players, comparing the number of identified equilibria $|\Phi|$, the runtime in seconds, and the conditional probability of the selected equilibrium across the different algorithms (SGM, eSGM-WM, eSGM-PO, eSGM). While the original SGM approach and the eSGM-PO are generally faster compared with the eSGM, they can only identify a subset of all equilibria. For the original SGM, the selected equilibrium solely depends on the initialization of the algorithm. Here, we initialize the SGM with the player's optimal strategies in the absence of competition. The eSGM-WM and the eSGM-PO, however, select the most likely solution, when players are restricted to the respective subset of equilibria (welfare-maximizing or Pareto-optimal equilibria, respectively). Note that this solution does not necessarily represent the most likely solution if players are not restricted to the respective subset of equilibria. If players can choose freely among all equilibria, the solution of the eSGM by far outperforms the extensions (eSGM-WM, eSGM-PO) as well as the SGM in terms of likelihood (Prob) of the identified equilibrium. In particular, there are some cases where the probability of the equilibrium selected by SGM, eSGM-WM and eSGM-PO are close or equal to zero, whereas the equilibrium determined by the eSGM has a probability greater than 50%. Figure 3.1a highlights the total number of mixed equilibria (denoted by $|\Phi_{\text{mixed}}|$) in the examined instances. Almost all equilibria are mixed. In 83% of all instances for $n = 2$ players, the most likely equilibrium is a mixed equilibrium. In most cases, this is due to the fact that no pure equilibrium exists, however there are instances where a mixed equilibrium is determined to be more likely although pure and mixed equilibria coexist. In Appendix 3.B, we provide an example and explanation for this observation.

Notably, the number of equilibria significantly decreases with the additional player in the case of $n = 3$, with a maximum of two equilibria per game for all instances ($|I| \in \{10, 20, 40\}$). With the reduced number of equilibria, the likelihood of the equilibrium identified through the SGM in most cases is equal to or closely resembles the likelihood of the equilibrium selected by the eSGM. Detailed results for the 3-player case can be found in the Appendix 3.E (Table 3.4).

3.3.2 Competitive facility location and design problem

Problem description

The competitive facility location and design problem combines location selection with design choices for each selected location. The problem was introduced by Aboolian et

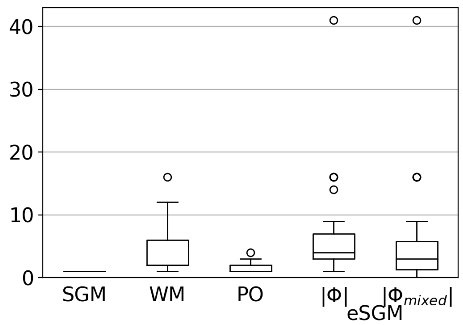

(a) Number of equilibria $|\Phi|$

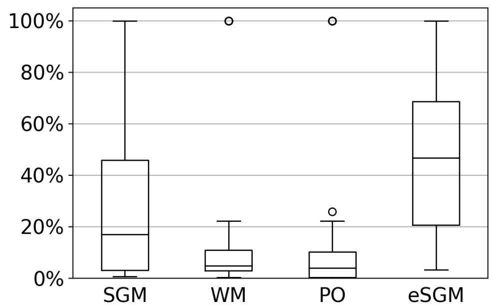

(b) Prob of the selected eq. among all eq. Φ

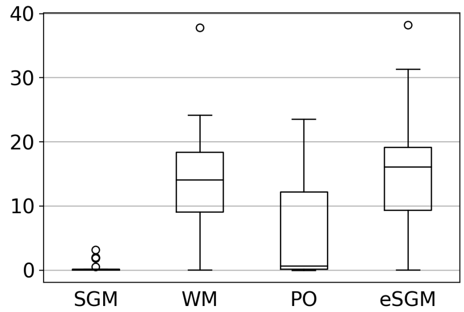

(c) CPU runtime in seconds

Figure 3.1: Knapsack game with $n = 2$, boxplots combine results for $|I| \in \{20, 40, 60\}$

al. (2007) to model market entry with known prior competitor locations. As part of the OR-Library, it has been adapted to a leader-follower (Stackelberg) setting, where a firm decides on the facility location and design under anticipation of its competitors' decisions (Beresnev, n.d.).

We present the following Nash problem formulation, where the competitors of set N decide simultaneously on the location and design of their facilities. Assuming a set I of potential facility locations and a set J of customers, R_i denotes the set of design alternatives for location $i \in I$. The fixed costs for the players $p \in N$ depend on location ($i \in I$) and design choices ($r \in R_i$) and are given by f^p_{ir}. The demand w_j of customer $j \in J$ is split between facilities based on the utility u^p_{ijr} a facility presents to the customer. We calculate u^p_{ijr} based on a design dependent location attractiveness a^p_{ir} and on the distance d_{ij} between the facility i and the customer j:

$$u^p_{ijr} = \begin{cases} \frac{a^p_{ir}}{(d_{ij}+1)^\beta} & , d_{ij} \leq d_{\max} \\ 0 & , d_{ij} > d_{\max} \end{cases} \tag{3.52}$$

Here, β denotes the distance sensitivity. In contrast to Beresnev (n.d.), we limit the distance up to which a customer can be served to $d_{\max}$. This reflects limited delivery ranges or a limited willingness of customers to travel long distances. Drezner et al. (2020) use a similar approach in their gradual cover model, where customer attraction is limited to a sphere of influence and, within this sphere of influence, gradually decreases with increasing distance to a facility.

Both players maximize the number of attracted customers:

$$\max_{x^p} \sum_{j \in J} w_j \frac{\sum_{i \in I} \sum_{r \in R_i} u^p_{ijr} x^p_{ir}}{\sum_{i \in I} \sum_{r \in R_i} u^p_{ijr} x^p_{ir} + \sum_{i \in I} \sum_{r \in R_i} u^{-p}_{ijr} x^{-p}_{ir}} \qquad \forall p \in N \tag{3.53}$$

with x^p_{ir} as the (binary) decision of player p to open a facility with design r at location i. The players are limited by their overall budget and can only choose one design per location:

$$\sum_{i \in I} \sum_{r \in R_i} f^p_{ir} x^p_{ir} \leq B^p \qquad \forall p \in N \tag{3.54}$$

$$\sum_{r \in R_i} x^p_{ir} \leq 1 \qquad \forall i \in I,\ p \in N \tag{3.55}$$

$$x^p_{ir} \in \{0, 1\} \qquad \forall i \in I,\ p \in N,\ \forall r \in R_i \tag{3.56}$$

The assumption of a limited sphere of influence of locations on customers ($d_{\max}$) implies that the examined competitive facility location and design problem is not a constant sum game, and social welfare ($\sum_{p \in P} \Pi^p$) can differ between equilibria.

Computational results

We use the benchmark data from Beresnev (n.d.), with $|I| = 50$ potential locations on a 100×100 euclidean square and $|J| = 50$ customers at the same locations. The companies A,B can choose between $|R_i| = 3$ different design alternatives for each location. Both companies share the same attractiveness ($a^p_{ir} = a^{-p}_{ir}$) and fixed costs ($f^p_{ir} = f^{-p}_{ir}$). Upon varying the budgets B^{A} and B^{B} and the parameters β and $d_{\max}$, we compare the runtime and solution quality of the proposed eSGM and its extensions with the SGM (Carvalho et al., 2022).

To highlight shortcomings of common normal-form algorithms, we also try to solve these problems using MIP-Nash (Sandholm et al., 2005). To do so, we first need to convert the IPG into normal-form through the explicit enumeration of all strategy combinations. Note that with a budget of $B^{\text{A}} = B^{\text{B}} = 40$, there are already more than 100mn possible strategy combinations. Hence, the required full enumeration of all strategy combinations exceeds the imposed time limit of 16 hours already at moderate problem sizes. As MIP-Nash is clearly outperformed by the sample-based approaches (SGM, eSGM and extensions) already for smaller instances, we do not report benchmark results in detail for MIP-Nash.

We test 4 different problem setups, varying the steepness (β) of the distance decay function and the distance cutoff $d_{\max}$. For each problem setup $(\beta, d_{\max})$ we solve 32 instances with different player budget combinations ranging from $B^{\text{A}} = B^{\text{B}} = 10$ to $B^{\text{A}} = 40, B^{\text{B}} = 80$. For all problems, the eSGM was able to solve all 32 instances within the 16 hour time limit. Table 3.1 shows an overview of the results, unless otherwise stated, we report averages across the 32 instances. The number of equilibria $|\Phi|$ varies greatly between problem setups. The majority of identified equilibria are mixed ($|\Phi_{\text{mixed}}|$), and in 53% of all instances, a mixed equilibrium is selected. While there are also several cases with a unique equilibrium where the SGM yields identical results in a shorter time, the SGM is not able to prove the uniqueness of this equilibrium. Most cases, however, exhibit multiple equilibria, leading to notable differences in solutions between the SGM and the eSGM, with the SGM often resulting in an equilibrium that is significantly less likely (based on Harsanyi's theory of player incentives) compared with the most probable equilibrium identified by the eSGM.

Further details, for the base case ($\beta = 0.5, d_{\max} = 20$) can be found in Table 3.2, with $|\Phi|, |\Phi_{\text{mixed}}|$ denoting the total number of equilibria and the number of mixed equilibria respectively, and with δ the competitiveness of each problem instance (see Definition 3.1). Although most equilibria are mixed, there are some instances with a unique equilibrium in pure strategies. In more homogeneous problem settings with similar player budgets, the competitiveness δ is higher compared to instances in which one player is predominant. If both players share the same budget ($B^{\text{A}} = B^{\text{B}}$), the competitiveness reaches its maximum observed value of 50%. The number of equilibria also increases with rising competitiveness δ, reaching up to 435 equilibria at maximum competitiveness ($\delta = 50\%$) in the case of $B^{\text{A}} = B^{\text{B}} = 20$. While the equilibrium identified through the SGM turns out to be the most probable or unique equilibrium in some cases, there are several cases where equilibrium identification through the SGM would significantly mislead decision makers. For example, in the case of $B^{\text{A}} = 10,\ B^{\text{B}} = 60$, the SGM returns an equilibrium with a probability of 3%, that is clearly outperformed by the most likely equilibrium (86%) identified by the eSGM.

Table 3.1: Computational results (32 instances per problem setup) for the competitive facility location and design problem

Problem setup		Number of equilibria		Runtime (1=SGM)			Prob⋆(%)	
β	$d_{\max}$	$\lvert\Phi_{\text{mixed}}\rvert_{\text{av}}/\lvert\Phi\rvert_{\text{av}}$	$\lvert\Phi\rvert_{\max}$	eSGM-WM	eSGM-PO	eSGM	SGM	eSGM
0.5	20	43/47	435	9.9	29.6	42.2	33	49
2	20	23/26	435	17.5	22.5	23.0	38	53
2	40	2/3	17	6.3	9.4	9.7	62	72
3	40	18/23	125	7.0	13.3	13.2	34	49

⋆ indicates the probability of the selected equilibrium among all equilibria.

3.4 Conclusions and future research

We introduce an integrated column- and row-generation method for the identification of all equilibria and the subsequent equilibrium selection in n-player finite mathematical programming games. This solution method enables us to identify the most likely solution. We show the application of the algorithm for knapsack games, large-scale competitive facility location (and design) problems and uncapacitated lot-sizing games of up to 3 players.

While the algorithm is capable of solving n-player games, the $(n-1)$-degree polynomial constraints make practical problem formulations for multiple ($n > 2$) players non-linear, regardless of the players' payoff functions. Appropriate decomposition or linearization methods could further reduce the computational effort for n-player games and should be investigated in further studies to enable practical application for $n > 3$.

The described solution method is explicitly targeted towards finite Mathematical programmming games (MPGs). For general MPGs, the original proof of equilibrium existence (Nash, 1951) would not hold. Even if the existence of equilibria could be guaranteed (see e.g., Stein et al. (2008) and Carvalho et al. (2018a) for mixed-integer programming games with separable payoff functions), the potentially infinite size of the set of candidate solutions could impede termination of the eSGM.

Potential for further research also lies in the relaxation of assumptions around behavioral rules. In particular, evolutionary games with discrete strategy sets or integer constraints (see e.g., Östling et al., 2011) are an interesting related area of research. The application to such a repeated game between boundedly rational players would require significant adaptations to the proposed method.

Table 3.2: Computational results in the base case ($\beta = 0.5,\ d_{\max} = 20$) for the competitive facility location and design problem

Player budget		2-player Characteristics		2-player Runtime (s)				2-player Prob⋆ (%)		3-player ($B^C = 10$) Equilibria	3-player Runtime (s)		3-player Prob⋆ (%)	
B^A	B^B	$\lvert\Phi_{\text{mixed}}\rvert/\lvert\Phi\rvert$	δ (%)	SGM	eSGM-WM	eSGM-PO	eSGM	SGM	eSGM	$\lvert\Phi_{\text{mixed}}\rvert/\lvert\Phi\rvert$	SGM	eSGM	SGM	eSGM
10	10	1/3	50	0.4	4.1	4.2	4.3	12	**75**	4/7	2.1	54.6	14	**18**
10	20	26/31	39	0.7	12.3	12.3	12.3	3	3	83/93	36.8	147	1	**1**
10	30	57/63	31	1.1	9.8	37.5	37.2	10	10	15/21	11.5	74.1	4	**6**
10	40	0/1	28	1	4.8	5.4	5.4	100	100	3/5	3.8	45.6	31	**38**
10	50	1/3	27	1.5	5.6	7.9	7.9	10	**84**	1/3	1.2	28.8	31	35
10	60	2/5	25	1.4	7.9	10	10.1	3	**86**	4/7	1.2	65.9	14	**18**
10	70	0/1	2	0.7	5.2	5.1	5.1	100	100	0/1	1	11.7	100	100
10	80	0/1	0	1.6	6.1	6.7	6.5	100	100	0/1	1.6	15.1	100	100
10	90	0/1	4	1.1	7.4	7.8	7.7	100	100	1/3	6.3	46.6	29	**41**
10	100	1/3	3	4	32	32.9	33.6	33	33	3/3	21	150	37	**37**
20	20	405/435	50	1.3	153.6	158.1	159.5	0	**1**	55/65	73.4	195.2	0	7
20	30	67/81	21	1.4	30.9	34.2	81.4	0	6	84/97	138.8	534	0	10
20	40	22/25	14	1.8	11.3	19.2	40.4	21	21	29/34	1.9	298.8	6	**10**
20	50	171/180	32	2.1	12.1	219.9	222.8	1	**3**	1/3	2	203.8	78	78
20	60	27/32	23	2.1	34.3	86.9	80.8	29	**40**	3/7	3.8	133.7	23	23
20	70	0/1	8	2.4	20.3	36.1	30.9	100	100	0/1	2.9	77.9	100	100
20	80	1/1	3	5.2	49.7	49.6	49.7	100	**100**	0/1	4.3	1260.2	100	100
20	90	1/1	7	24.5	136.4	135.9	136.3	100	**100**	3/3	15	2534.5	33	**33**
20	100	3/3	3	62	2850.4	2845.6	2836.6	0	**33**	3/3	97	6430.6	33	**33**
30	30	277/283	50	4.9	24.8	578.8	689.6	2	**6**	47/59	48.4	670.3	2	**11**
30	40	3/5	40	5	10.1	55.9	277.1	20	**58**	31/40	273.9	2335.6	0	**10**
30	50	24/31	33	3.1	85.6	502.8	502.6	0	**11**	37/43	17.9	1387.7	14	**20**
30	60	4/9	26	4.6	1319.4	650.8	497.4	3	**63**	4/7	15.1	2958	7	**23**
30	70	3/5	16	1.9	86.9	508.7	511.8	73	73	4/5	4.1	2339	69	69
30	80	4/5	9	8.4	9994.5	6861.2	3644.9	1	**32**	12/13	15.3	23517.2	77	77
30	90	15/15	7	20.3	6445.2	6525.3	6530	0	**10**	3/3	75.6	14403.8	33	**33**
30	100	18/18	2	14661.5	20531.5	22234.9	22201.5	0	**8**	1/3	90.3	35630.2	33	33
40	40	187/197	46	6.6	483.7	20473.4	5767.2	0	**1**	463/473	3409.6	16990.6	0	**0**
40	50	36/41	38	8.5	708	2480.8	2910.9	0	**26**	174/185	33.2	20461.1	1	**3**
40	60	7/11	25	7.1	6322.1	5041.2	3638.8	1	**53**	30/35	15.7	32635	3	7
40	70	3/5	12	6.2	19647.9	5783.6	7102.1	44	44	15/21	11.2	38507.7	4	44
40	80	0/1	9	11.2	138362.6	34607.3	49779.1	100	100	14/19	73.4	95888.1	9	10

⋆ indicates the probability of the selected equilibrium among all equilibria. **Bold** probabilities indicate that the equilibrium selected through the eSGM (highest probability) is a mixed equilibrium.

Appendix 3.A Monte Carlo volume estimation of W_{x^p}

As the mapping μ^p transforming V_{x^p} into W_{x^p} is non-linear, in contrast to V_{x^p}, the set W_{x^p} has non-linear bounds. To efficiently compute the volume of W_{x^p} we apply a numerical estimation approach to determine the size $\lambda(W_{x^p})$. We use the simple and widely applied hit-or-miss method for Monte Carlo volume estimation as presented in Fok and Crevier (1989). Just as V_{x^p}, the set W_{x^p} is a subset of a $(|S_\mathrm{r}^{-p}|-1)$-dimensional simplex. For W_{x^p}, we can define this simplex Δ as:

$$\Delta = \left\{ r^p \in \mathbb{R}^{|S_\mathrm{r}^{-p}|} \,\middle|\, r^p_{x^{-p}} \geq 0, \sum_{x^{-p} \in S_\mathrm{r}^{-p}} r^p_{x^{-p}} = 1 \right\} \tag{3.57}$$

We sample s points $\tilde{r}^p \sim \mathrm{Dir}(\boldsymbol{\alpha})$ from a Dirichlet distribution, where the parameter vector $\boldsymbol{\alpha}$ consists of all ones, i.e., a uniform distribution over this simplex. Let $s_{W_{x^p}}$ be the number of sampled points that lie within W_{x^p} (i.e., $\tilde{r}^p \in W_{x^p}$). Given a large enough sample size s, the size $\lambda(W_{x^p})$ can be approximated as:

$$\lambda(W_{x^p}) \approx \frac{s_{W_{x^p}}}{s} \lambda(\Delta) \tag{3.58}$$

Note that although the size $\lambda(\Delta)$ of the simplex Δ can be easily calculated, we do not require any information about $\lambda(\Delta)$ as we determine a player's probability of a strategy profile based on the relative incentive of the strategy profile over all alternatives (3.37). In our implementation we use a sample size of $n = 1 \times 10^5$, yielding consistent results for the dimensionality encountered in the examined numerical studies.

Appendix 3.B Equilibrium selection in the knapsack game

Let us assume a setting for the knapsack game (Section 3.3.1), where two players $N = \{\mathrm{A}, \mathrm{B}\}$ select one out of $|I| = 2$ technologies. Imagine that a given technology $i \in I$ will only pay off if both competitors select it, meaning that $v_i^p = 0$. In this case, if both competitors prefer the same technology, equilibrium selection would be trivial. In the following, we assume A strongly prefers $i = 1$, and B strongly prefers $i = 2$:

In a symmetric case where $c_1^{\mathrm{AB}} = 1$, $c_2^{\mathrm{AB}} = 10$ and $c_2^{\mathrm{BA}} = 10$, $c_2^{\mathrm{BA}} = 1$, we have two pure Nash equilibria (with both players successfully jointly selecting either technology) and a mixed equilibrium with combinations of the two technologies, yielding the set of all equilibria: $\Phi = \left\{((1,0),(1,0)),((0,1),(0,1)),\left(\left(\frac{1}{11},\frac{10}{11}\right),\left(\frac{10}{11},\frac{1}{11}\right)\right)\right\}$.

Note as the players only have two pure strategies each, the stability set $V_{x^{\mathrm{A}}}$ of player A using strategy x^{A} is a line. Let $x^{\mathrm{A}}_{i=1} = (1,0)$ be a pure strategy where A selects technology 1 and $x^{\mathrm{A}}_{i=2} = (0,1)$ the alternate pure strategy where A selects 2. Figure 3.2 shows the stability sets $V_{x^{\mathrm{A}}_{i=1}}$ and $V_{x^{\mathrm{A}}_{i=2}}$ for player A.

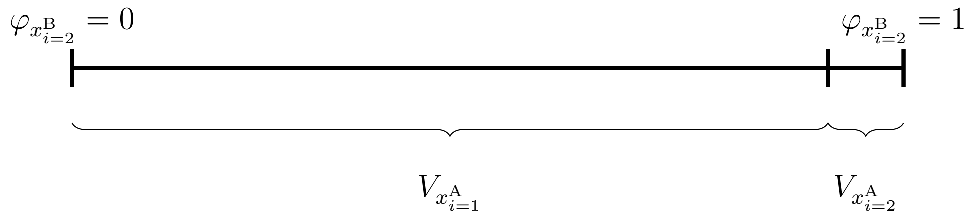

Figure 3.2: Stability sets $V_{x^{\mathrm{A}}_{i=1}}$ and $V_{x^{\mathrm{A}}_{i=2}}$ for the coordination knapsack game

Using the function μ^p, we determine the respective sets $W_{x^{\mathrm{A}}}$ and calculate the incentive $\lambda(W_{x^{\mathrm{A}}})$ for A to play a strategy x^{A}. Following the same process for B, we arrive at the following incentives for pure strategies:

$$\lambda(W_{x^{\mathrm{A}}_{i=1}}) = \lambda(W_{x^{\mathrm{B}}_{i=2}}) = \frac{1}{11}$$
$$\lambda(W_{x^{\mathrm{A}}_{i=2}}) = \lambda(W_{x^{\mathrm{B}}_{i=1}}) = \frac{10}{11}$$

For the mixed equilibrium $\hat{\varphi} = \left(\left(\frac{1}{11},\frac{10}{11}\right),\left(\frac{10}{11},\frac{1}{11}\right)\right)$ we get:

$$\Psi^{\mathrm{A}}_{\hat{\varphi}^{\mathrm{A}}} = \frac{1}{11}\lambda(W_{x^{\mathrm{A}}_{i=1}}) + \frac{10}{11}\lambda(W_{x^{\mathrm{A}}_{i=2}}) = \frac{101}{121}$$
$$\Psi^{\mathrm{B}}_{\hat{\varphi}^{\mathrm{B}}} = \frac{10}{11}\lambda(W_{x^{\mathrm{B}}_{i=1}}) + \frac{1}{11}\lambda(W_{x^{\mathrm{B}}_{i=2}}) = \frac{101}{121}$$

Scaling Ψ to ensure a cumulative probability of one per player, yields the probabilities:

$$\mathrm{Prob}(\Theta^{\mathrm{A}} = x^{\mathrm{A}}_{i=1}) = \mathrm{Prob}(\Theta^{\mathrm{B}} = x^{\mathrm{B}}_{i=2}) = 5\%$$
$$\mathrm{Prob}(\Theta^{\mathrm{A}} = x^{\mathrm{A}}_{i=2}) = \mathrm{Prob}(\Theta^{\mathrm{B}} = x^{\mathrm{B}}_{i=1}) = 50\%$$
$$\mathrm{Prob}(\Theta^{\mathrm{A}} = \hat{\varphi}^{\mathrm{A}}) = \mathrm{Prob}(\Theta^{\mathrm{B}} = \hat{\varphi}^{\mathrm{B}}) = 45\%$$

This suggests that successful coordination to either of the two pure equilibria is equally unlikely with $\mathrm{Prob}(\Theta^{\mathrm{A}} = x^{\mathrm{A}}_{i=1}) \cdot \mathrm{Prob}(\Theta^{\mathrm{B}} = x^{\mathrm{B}}_{i=1}) = \mathrm{Prob}(\Theta^{\mathrm{A}} = x^{\mathrm{A}}_{i=2}) \cdot \mathrm{Prob}(\Theta^{\mathrm{B}} =$

$x^{\mathrm{B}}_{i=2}) = 2.25\%$, whereas the mixed equilibrium $\hat{\varphi}$ is the expected outcome of the game with $\mathrm{Prob}(\Theta^{\mathrm{A}} = \hat{\varphi}^{\mathrm{A}}) \cdot \mathrm{Prob}(\Theta^{\mathrm{B}} = \hat{\varphi}^{\mathrm{B}}) = 24.55\%$.

Appendix 3.C Process flow diagram

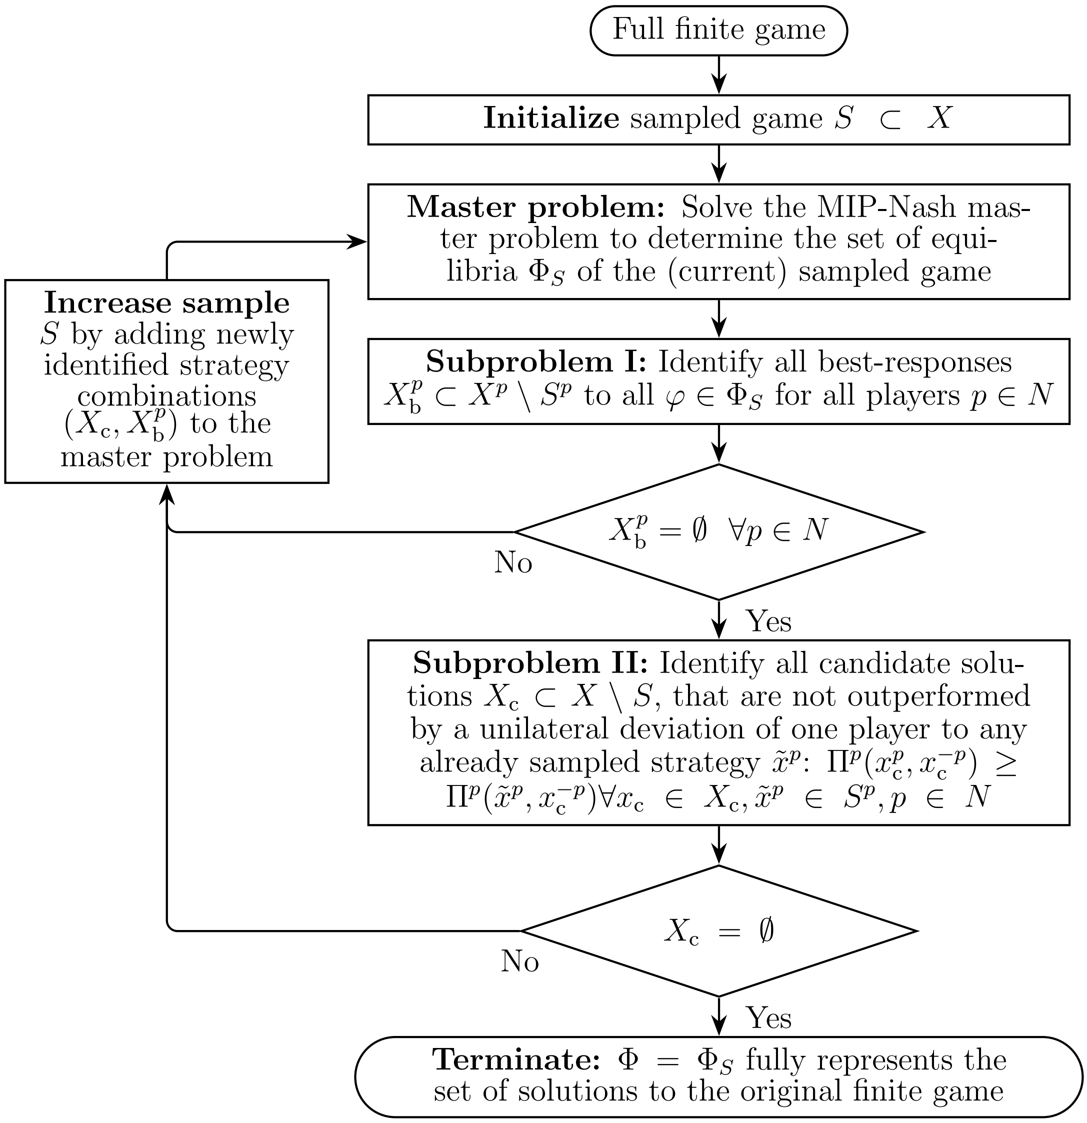

Figure 3.3: Process flow diagram of the proposed eSGM

Appendix 3.D Combinatorial competitive uncapacitated lot-sizing game

3.D.1 Problem description

Carvalho et al. (2018b) present a combinatorial variant of the competitive uncapacitated lot-sizing problem. In this combinatorial problem, n players produce a (single) homogeneous good in a finite planning horizon $|T|$. The price of the good $p_t(q_t)$ is determined by the market based on the total quantity $q_t = \sum_{p\in N} q_t^p$ of offered units across all players: $p_t(q_t) = \max\{a_t - b_t q_t, 0\}$, with $a_t, b_t \geq 0$ as market characteristics and q_t^p as the quantity offered by player p in period t. z_t^p denotes production quantities, the binary variable $y_t^p \in \{0,1\}$ denotes production periods. In each period of active production ($y_t^p = 1$), fixed setup costs f^p are incurred. Goods that are produced but not sold $z_t^p > q_t^p$ increase the inventory h_t^p. As Carvalho et al. (2018b), we assume no holding costs and focus on the combinatorial decision problem of finding optimal periods of active production y_t^p. Decisions on sales and production quantities (q_t^p and z_t^p, respectively) follow from the zero-inventory property (Wagner and Whitin, 1958). Players optimize individual profits Π^p as the difference between operational margin and fixed costs, with $\bar{u}_t^p$ defining the last period of active production prior to and including t and $c^p_{\bar{u}_t^p}$ as the production cost in this period (Carvalho et al., 2018b):

$$\max_{y_t^p, y_t^{-p}} \Pi^p(y_t^p, y_t^{-p}) = \sum_{t\in T | \sum_{\tau\in\{0,\ldots,t\}} y_\tau^p \geq 1} \bar{q}_t^p (p_t(q_t) - c^p_{\bar{u}_t^p}) - \sum_{t\in T} f^p y_t^p \tag{3.59}$$

In case a period is not preceded by a period of active production, $\bar{u}_t^p$ and by consequence $c^p_{\bar{u}_t^p}$ are undefined. Therefore, in (3.59), we limit the sum of the operational margin to periods with active production or periods that are preceded by a period of active production.

$$\bar{u}_t^p = \max\{u \mid u \in T, u \leq t, y_u^p = 1\} \qquad \forall t \in T,\ \forall p \in N \tag{3.60}$$

In (3.60), $\bar{u}_t^p$ is formally defined as the period $u \in T$ with active production ($y_u^p = 1$) and the highest cardinality ($\max u$) smaller or equal to t.

Note that (3.59) expects a decision only on y_t^p, y^{-p}. To translate the decision on the set of production periods into optimal production quantities $\bar{q}_t^p$, Carvalho et al. (2018b) derive an optimality condition by setting the partial derivative of Π^p with regards to $\bar{q}_t^p$

to zero:

$$\frac{\partial \Pi^p}{\partial \bar{q}_t^p} = 0 \rightarrow \bar{q}_t^p = \frac{\max\left\{a_t - b_t \sum_{\tilde{p} \neq p} q_t^{\tilde{p}} - c^p_{\bar{u}_t^p}, 0\right\}}{2b_t} \qquad \forall p \in N \tag{3.61}$$

Reformulation yields:

$$\bar{q}_t^p = \begin{cases} \frac{\max\left\{p_t - c^p_{\bar{u}_t^p}, 0\right\}}{b_t} & \text{, if } \sum_{\tau \in \{0,\ldots,t\}} y_\tau^p \geq 1 \\ 0 & \text{, else} \end{cases} \qquad \forall t \in T,\ \forall p \in N \tag{3.62}$$

$$p_t = \frac{a_t + \sum_{p \in N | y_t^p = 1} c^p_{\bar{u}_t^p}}{\sum_{p \in N} y_t^p + 1} \qquad \forall t \in T \tag{3.63}$$

This combinatorial model features no independent continuous decision variables. We solve the model using eSGM with the following implementation of (3.59) and (3.60): First, we replace $\max\{p_t - c^p_{\bar{u}_t^p}, 0\}$ with a new auxiliary variable k_t^p.

$$\Pi^p(y_t^p, y_t^{-p}) = \sum_{t \in T} \frac{(k_t^p)^2}{b_t} - \sum_{t \in T} f^p y_t^p \tag{3.64}$$

We introduce $o_t^p \in \{0,1\}$ as an auxiliary binary variable and ensure $k_t^p = \max\{p_t - C^p_{\bar{u}_t^p}, 0\}$ as follows:

$$k_t^p \leq \sum_{\tau \in \{0,\ldots,t\}} M y_\tau^p \qquad \forall t \in T,\ \forall p \in N \tag{3.65}$$

$$k_t^p \leq p_t - c^p_{\bar{u}_t^p} + o_t^p M \qquad \forall t \in T,\ \forall p \in N \tag{3.66}$$

$$k_t^p \geq M(o_t^p - 1) \qquad \forall t \in T,\ \forall p \in N \tag{3.67}$$

$$M(1 - o_t^p) \geq p_t - c^p_{\bar{u}_t^p} \qquad \forall t \in T,\ \forall p \in N \tag{3.68}$$

Whenever $p_t - c^p_{\bar{u}_t^p} \leq 0$, (3.68) sets $o_t^p = 1$ and by consequence $k_t^p = 0$ (3.67). In all other cases k_t^p is maximized according to the objective function (3.64) and hence equal to $k_p^t = p_t - c^p_{\bar{u}_t^p}$ according to (3.66). Furthermore, to calculate $c^p_{\bar{u}_t^p}$, we introduce the binary auxiliary variable $u_{t\tau}^p \in \{0,1\}$. We ensure $u_{t\tau}^p = 1$, if and only if τ was the last period of active production for p prior to t (i.e., $\bar{u}_t^p = \tau$):

$$u_{t\tau}^p = 0 \qquad \forall t \in T,\ \tau \in \{t+1, \ldots, |T|\},\ \forall p \in N \tag{3.69}$$

$$u_{t\tau}^p \leq y_\tau^p \qquad \forall t, \tau \in T,\ \forall p \in N \tag{3.70}$$

$$\sum_{\tau \in T} u^p_{t\tau} \leq 1 \qquad \forall t \in T,\ \forall p \in N \tag{3.71}$$

$$u^p_{t\tau} + (1 - y^p_\tau)M \geq u^p_{t\hat{t}} \qquad \forall t \in T,\ \forall \tau \in \{0, \ldots, t\},\ \forall \hat{t} \in \{0, \ldots, \tau\},\ \forall p \in N \tag{3.72}$$

$$u^p_{tt} = y^p_t \qquad \forall t \in T,\ \forall p \in N \tag{3.73}$$

(3.69) sets $u_{t\tau}$ to zero for all periods $\tau > t$. (3.70) ensures that $u_{t\tau}$ can only be one, if τ is a period of active production. As we determine the (single) last period of active production, the sum of $u_{t\tau}$ over $t \in T$ needs to add up to one (3.71), and a more recent active period always needs to supersede a later period of active production (3.72). Using $u^p_{t\tau}$, we can calculate $c^p_{\bar{u}^p_t}$ as:

$$c^p_{\bar{u}^p_t} = c^p_t \sum_{\tau \in T} u^p_{t\tau} \qquad \forall t \in T,\ \forall p \in N \tag{3.74}$$

The combinatorial problem (3.59)-(3.74) builds on the property that players decide on the quantity $\bar{q}^p_t$ according to optimality condition (3.61). If only a single player was to unilaterally change their production periods (and by consequence quantity), all other players are expected to implicitly adapt their production $\bar{q}^p_t$ to their respective new optimum. Carvalho et al. (2018b) show that this implies that the set of equilibria of the combinatorial problem is a subset of the equilibria of an alternative mixed integer formulation, as equilibria in the combinatorial problem are not only stable against unilateral changes but also against (implicit) reactions in production quantities from competitors.

3.D.2 Computational results for the competitive uncapacitated lot-sizing game

We draw the parameters a_t and b_t independently from discrete uniform distributions for three different time horizons $|T| \in \{10, 20, 50\}$: $a_t \sim \mathcal{U}\{20, 30\},\ b_t \sim \mathcal{U}\{3, 5\}$. For each time horizon, we generate 10 instances for $n = 2$ and $n = 3$ players. We assume setup costs are constant $f^p = 1$, whereas variable costs start with $\bar{c}^p \sim \mathcal{U}\{29, 31\}$ for $t = 0$ and decrease gradually over the time horizon T: $c^p_t = \bar{c}^p \left(1 - 0.35\frac{t}{|T|}\right)$. The motivation for this type of cost function is twofold, as it allows to model learning behavior or continuous performance improvements and provides an indirect way to account for holdings costs: Producing goods in period $t-1$ to sell in period t leads to an increase in cost by factor $\frac{0.35}{|T|}$.

Table 3.3 shows aggregate (min, mean, max) results for the runtime (RT) and the identified number of equilibria $|\Phi|$ for each combination $(n, |T|)$ using SGM and eSGM. Using the eSGM, we show that for the majority of examined instances there is no unique equilibrium, and some instances show up to 70 equilibria. More than a third of all equilibria are mixed, but only in 7% of instances the selected equilibrium is a mixed one. While on average, the conditional probability of the equilibrium selected by the eSGM only marginally outperforms the unique equilibrium identified using SGM, there are some cases where the probability of the latter is almost equal to zero whereas the eSGM selects a highly probable (> 50%) equilibrium.

Table 3.3: Computational results for the competitive uncapacitated lot-sizing game (RT=runtime in s)

Problem		SGM				eSGM						
n	$\|T\|$	RT_{min}	RT_{mean}	RT_{max}	Prob⋆(%)	RT_{min}	RT_{mean}	RT_{max}	$\|\Phi\|_{\text{av}}$	$\|\Phi_{\text{mixed}}\|_{\text{av}}$	$\|\Phi\|_{\text{max}}$	Prob⋆(%)
2	10	0.0	0.1	0.2	93	0.1	0.3	0.7	1.8	0.4	7	93
	20	0.3	0.5	0.8	82	0.6	3.6	15.1	3.8	1.8	17	84
	50	3.5	3.9	4.6	76	7.7	172.2	1533.9	4	1.8	17	83
3	10	0.1	0.1	0.3	91	0.2	0.7	2.2	2.1	0.7	10	93
	20	0.4	0.6	1.2	93	2.4	46.6	372.8	2	0.5	5	93
	50	2.3	5.1	8.0	62	20.7	1772.5	15968.2	9.7	7.4	70	77

⋆ indicates the probability (in %) of the selected equilibrium among all equilibria.

Appendix 3.E Knapsack game: 3-player results

Table 3.4: Computational results for the coordination knapsack game with $n = 3$

Instances		SGM			eSGM			
$\|I\|$	m	$\|\Phi\|$	time (s)	Prob*(%)	$\|\Phi\|$	$\|\Phi_{\text{mixed}}\|$	time (s)	Prob*(%)
10	0	1	0	100	1	0	0.1	100
	1	1	0	100	1	0	0.1	100
	2	1	0	100	1	0	0.1	100
	3	1	0	100	1	0	0.1	100
	4	1	0	100	1	0	0.2	100
	5	1	0.5	71	2	2	15.1	**71**
	6	1	0.1	100	1	0	0.3	100
	7	1	1.4	29	2	1	13.9	71
	8	1	0.1	100	1	0	0.2	100
	9	1	0.1	100	1	0	0.2	100
20	0	1	0	100	1	0	0.1	100
	1	1	0.1	100	1	0	0.2	100
	2	1	0.1	100	1	0	0.3	100
	3	1	0	100	1	0	0.5	100
	4	1	0.1	100	1	0	1.0	100
	5	1	0.1	100	1	0	1.6	100
	6	1	0.1	100	1	0	1.1	100
	7	1	0.1	100	1	0	0.5	100
	8	1	0.1	100	1	0	0.5	100
	9	1	0.1	100	1	0	0.5	100
40	0	1	0	100	1	0	0.2	100
	1	1	1.3	50	2	2	10.3	**50**
	2	1	7.6	21	2	2	43.3	**79**
	3	1	0.1	100	1	0	10.5	100
	4	1	0.2	99	2	0	38.6	99
	5	1	0.2	2	2	0	52.3	98
	6	1	0.1	100	1	0	53.7	100
	7	1	0.1	100	1	0	55.1	100
	8	1	0.1	100	1	0	48.9	100
	9	1	0.1	100	1	0	55.5	100

$\star$ indicates the probability of the selected equilibrium among all equilibria. **Bold** probabilities indicate that the equilibrium selected through the eSGM (highest probability) is a mixed equilibrium.

Chapter 4

Location selection for hydrogen fuel stations under emerging provider competition

Individual hydrogen-based mobility with Fuel cell electric vehicles (FCEVs) is a promising avenue for green house gas reduction in the transportation sector. With a rise in popularity of FCEVs and increased governmental grants and funding, investments in hydrogen refueling stations gain attractiveness for fuel station providers. However, as uncertainty around adoption rates and emerging competition between station providers persists, careful location selection becomes crucial for providers willing to invest in hydrogen fuel stations.

We formulate a location problem for hydrogen fuel stations as a Competitive flow capturing location-allocation model (C-FCLM). In contrast to prior formulations, no prior knowledge of competitor locations is required as we expect competitors to choose locations simultaneously in an emerging competitive environment. We solve the arising competitive model formulation as an Integer programming game (IPG). In a real-world scenario, we identify optimal locations for the two largest competing fuel station providers in Munich, Germany. Within this study, providers who acknowledge the existence of their competitor can realize an increase in profit of 17% (averaged across different scenarios of FCEV market penetration & customer preferences). Central coordination (e.g., through conditional subsidies from policymakers) or cooperation between the competitors (e.g., through provider associations) could further increase overall profits by 28%, however this increase in profitability comes at the cost of overall customer travel distance (increased by up to 10% in case of cooperation between providers) as the average detour a customer has to take to refuel increases.

4.1 Introduction

4.1.1 Motivation

Compared with other sectors (e.g., energy, industry or agriculture), the transport sector significantly lacks behind in green house gas reduction, showing a net increase in greenhouse gases (CO2 equivalent) of 23% compared to 1990 on a European level (European Environment Agency, 2020). Alternative fuels constitute an important lever to tackle this gap (McKinnon et al., 2015, pp. 278). To achieve a significant reduction in green house gas levels, a multitude of policies favoring alternative fuels have been put in place in individual states as well as internationally in the recently published European Hydrogen Strategy (European Commission, 2020). In addition to these policies, major technological breakthroughs (e.g., the liquefaction of hydrogen in organic carriers, reductions in battery production costs or the efficient synthetization of e-fuels) promise significant contributions to green house gas reduction in transportation.

However, these major changes bring along significant uncertainties for decision makers: The adoption of new propulsion technologies by customers is still largely unclear (Apostolou and Xydis, 2019; Lyons and Davidson, 2016) and the nature of newly developing markets (e.g., hydrogen-based mobility) implies an unclear competitive landscape and uncertainty around competitor actions. Among the main entry barriers limiting adoption of hydrogen mobility remain the lack of available Hydrogen fuel stations (HFSs) (Ramea, 2019) and the high construction costs of these stations (Viktorsson et al., 2017).

Service providers trying to offer HFSs thus face two major issues: In comparison with Electric vehicle (EV) charging stations, investment costs for hydrogen fueling stations are higher (i). This makes location decisions inflexible and planning errors costly. In contrast to conventional fuel stations, there is little prior knowledge about competitor locations for HFSs, as multiple providers might scale up hydrogen fueling station networks almost simultaneously (ii). Location decisions that are optimal under negligence of possible competitor decisions, could thus turn out sub-optimal or even unprofitable under consideration of imminent competitor actions. This second issue especially applies to location decisions in urban areas where a dense coverage (and thus competition) of fuel stations can be expected. It might not be as apparent for fuel stations on highways, with only 2.4% of all (conventional) fuel stations in Germany being located on a highway (Köhler et al., 2010).

In the absence of a pre-existing hydrogen fuel station infrastructure, we cannot rely on a Stackelberg-type formulation with a clear leader and a clear follower. Instead, it is to be expected that fuel station providers will enter this newly emerging market almost simultaneously. Early movers can create effective entry barriers through the preemption of scarce assets (Lieberman and Montgomery, 1988). For HFSs this early mover advantage is particular relevant based on the limited market size in the near-term: Once a critical mass of hydrogen fuel stations is reached, late entries cannot reach the market share required for profitability. To avoid this entry barrier, providers cannot follow a wait-and-see approach but instead need to make decisions now while taking into account possible (simultaneous) decisions by their competitors. While there might be followers in case hydrogen mobility adoption rises in the long-term, the near-term competitive environment is best modeled using a Nash setting of simultaneous competition between early movers.

4.1.2 Research question and contribution

The main objective of our research is two-fold: First, we model the inner-city location decision for service providers building HFSs under consideration of competitor actions. Second, we apply the resulting problem formulation to a case study in Munich to demonstrate computational tractability for large scale applications and to deduct managerial insights. The resulting research questions (RQ) can be detailed as follows:

RQ 2.1 *How will competition between station providers influence the emerging hydrogen refueling network structure?*

RQ 2.2 *How valuable is it for decision makers to take competitor actions into consideration?*

RQ 2.3 *Should policymakers foster (e.g., through government-backed provider associations) or impede (e.g., through strict antitrust laws) collaboration between competing providers?*

We extend prior models for alternative fuel station location selection to a simultaneous competitive setting. Hereby we present the C-FCLM, which is not reliant on a priori knowledge of competitor locations. We establish analogies to well researched congestion games, to determine the existence of Nash equilibria for this model. As Flow capturing location-allocation models (FCLMs) rely on granular data of individual trip origins and destinations, we quantify the impact of trip aggregation in case granular data is not available or results in intractable problem sizes.

4.1.3 Organization

Section 4.2 presents the general C-FCLM. Section 4.3 introduces a solution algorithm, characteristics of the introduced model and solution aspects. We conduct a numerical study in Section 4.4 and showcase practical application for hydrogen fueling station location selection in Munich in Section 4.5. We conclude results and further research opportunities in Section 4.6.

4.2 The competitive flow capturing location-allocation model (C-FCLM)

We extend the classical flow-capture location model to a competitive scenario between two or more decision makers (C-FCLM). When developing and applying the C-FCLM, we assume facilities to be uncapacitated. Competitors are assumed to be purely self-interested and have full knowledge of potential competitor actions (strategies) and their cost structure. In our application of the C-FCLM, we focus on an inner-city setting. Given the high driving range of FCEVs, we can therefore neglect any range limitations. We assume that all customers sharing the same origin and destination intend to follow the same (shortest) path, neglecting potential congestion effects. Customer patronage of refueling facilities is assumed to be driven by the detour (compared with the originally intended shortest path) a customer has to take to arrive at the facility.

4.2.1 General C-FCLM

We examine a road network of $|K|$ connected nodes indexed with $i, j \in K$. The traffic flow through the network is given by flow volume f_q, representing the number of cars or potential customers per year on route $q \in Q$. We represent q by its origin-destination (OD) pair $q = (i, j)$, with both origin $i \in K$ and destination $j \in K$ representing nodes of the original network. Consider a set P of competing fuel station providers. Each provider $p \in P$ wants to find optimal locations within the set of potential locations $K_p \subset K$, maximizing individual profit Π_p (4.1). We calculate the profit as the difference between operating margin and annual fix costs (including depreciation and maintenance). The operating margin is driven by the fraction of flow on q that p can capture ($y_{qp} \in [0, 1]$), the number of customers f_q on q and the margin per customer m. Fix costs are determined by the binary decision to open a location $x_{kp} \in \{0, 1\}$ and the annual cost per facility c.

$$\max_{x_{kp}} \Pi_p = m \sum_{q \in Q} f_q y_{qp} - c \sum_{k \in K_p} x_{kp}, \tag{4.1}$$

Players are restricted to their respective set of potential locations K_p:

$$x_{kp} = 0 \qquad \forall k \in K \setminus K_p, \ \forall p \in P \tag{4.2}$$

Each customer can only be captured once, no matter how many active nodes it passes. By consequence, the cumulatively (across all players) captured fraction y_{qp} of the flow f_q on one particular OD-pair q cannot exceed one:

$$\sum_{p \in P} y_{qp} \leq 1 \qquad \forall q \in Q \tag{4.3}$$

The fraction y_{qp} of customers on OD pair q that player p can capture is driven by assumptions on customer choice behavior: In their basic version, market capture models for facility location allocate all customers of a given demand node to the single facility with the highest customer proximity (Hotelling, 1929; Serra et al., 1999). Relaxing this all-or-nothing property, other market capture models typically distribute customers on facilities relatively to a measure of attractiveness, most notably the MAXPROP formulation by Serra et al. (1999), and the gravity model of Huff (1964). Building on discrete choice analysis (see Ben-Akiva and Lerman, 1997; McFadden, 1974), the application of Multinomial logit choice (MNL) models to estimate market capture based on deterministic (observable) and random (non-observable) utility gained popularity in various applications (e.g., Benati and Hansen, 2002; Haase and Müller, 2013; Marianov et al., 2008; Zhang et al., 2012). We opt to extend the application to HFSs and model y_{qp} based on a MNL model. We assume that the total utility of a facility k for customers on OD-pair q is given by:

$$\tilde{u}_{qk} = u_{qk} + \tilde{\epsilon}_{qk} \qquad \forall q \in Q, \ \forall k \in K \tag{4.4}$$

$\tilde{\epsilon}_{qk}$ represents the random, non-observable part of the utility, which we assume to be independently and identically Gumbel distributed. We determine the deterministic utility u_{qk} of a facility k to customers of OD-pair q based on the facility attractiveness a_k (e.g., based on facility size or brand) and the deviation distance d_{qk} a customer faces when patronizing facility k. We calculate d_{qk} as the difference between the shortest path length that connects origin and destination of q through the facility k and the length of the original (shortest) OD-path:

$$d_{qk} = d(i,k) + d(k,j) - d(i,j) \qquad \forall q = (i,j) \in Q,\ k \in K \tag{4.5}$$

Here, $d(i,j)$ is the shortest path distance between nodes $i,j \in K$. Assuming that a customer on $q = (i,j)$ always takes the shortest path between origin i and destination j, $d(i,k) + d(k,j) > d(i,j)$ and hence $d_{qk} \geq 0$. As in the FCLM customers are assumed to refuel along an originally planned route rather than a single purpose trip, the customers' willingness to take long detours to refuel is limited. We impose a limit of $\bar{d}_q$, which is calculated based on the minimum of an absolute upper bound $\bar{d}$ and a relative willingness to travel $\hat{d}$ multiplied with the driving distance of the originally planned trip $d(i,j)$:

$$\bar{d}_q = \min\left(\hat{d} \cdot d(i,j),\ \bar{d}\right) \qquad \forall q = (i,j) \in Q \tag{4.6}$$

Given a deviation distance d_{qk} and a facility attractiveness a_k, we calculate u_{qk} as:

$$u_{qk} = \begin{cases} a_k^{\alpha} \frac{(\bar{d}_q - d_{qk})^{\beta}}{\bar{d}_q^{\beta}} & \text{for } \bar{d}_q \geq d_{qk} \\ 0 & \text{for } \bar{d}_q < d_{qk} \end{cases} \tag{4.7}$$

where $\alpha \geq 0$ and $\beta \geq 0$ are parameters reflecting customer sensitivities for a_k and d_{qk}, respectively. The utility increases with a higher attractiveness a_k and decreases with longer deviations d_{qk}, it reaches zero when the deviation distance equals the maximum willingness to travel $\bar{d}_q$. Instead of penalizing facilities located outside of this willingness to travel ($d_{qk} > \bar{d}_q$) with a negative utility, we fully excluded these facilities from the choice set for customers on the respective OD-trip q. As we model an inner-city setting, there is no obligation to refuel on this specific trip, and no rational customer would choose a facility outside of their own willingness to travel over the no-choice alternative. Hence, the probability of a customer q choosing an (active) facility k is (Ben-Akiva and Lerman, 1997):

$$y_{qk} = \frac{e^{u_{qk}}}{\sum\limits_{\tilde{k} \in K | d_{q\tilde{k}} \leq \bar{d}_q} e^{u_{q\tilde{k}}} \sum_{\tilde{p} \in P} x_{\tilde{k}\tilde{p}}} \qquad \forall k \in K,\ \forall q \in Q \ \middle|\ \bar{d}_q \geq d_{qk},\ \sum_{\tilde{p} \in P} x_{k\tilde{p}} > 0 \tag{4.8}$$

This approach enables a sparse representation, where $y_{qk} = 0$ for all locations outside of the customers reach ($\forall k \in K,\ \forall q \in Q \mid d_{qk} > \bar{d}_q$). The overall market share of q for player p is the cumulative probability across all active facilities that are able to serve q:

$$y_{qp} = \sum_{k \in K_p | d_{qk} \leq \bar{d}_q} y_{qk} x_{kp} = \frac{\sum\limits_{k \in K_p | d_{qk} \leq \bar{d}_q} e^{u_{qk}} x_{kp}}{\sum\limits_{\tilde{k} \in K | d_{q\tilde{k}} \leq \bar{d}_q} e^{u_{q\tilde{k}}} \sum_{\tilde{p} \in P} x_{\tilde{k}\tilde{p}}} \qquad \forall p \in P,\ \forall q \in Q \tag{4.9}$$

We linearize (4.9) by introducing the auxiliary variables $y_{\tilde{k}\tilde{p}qp} \in [0, 1]$, with the following additional constraints ensuring that $y_{\tilde{k}\tilde{p}qp} = x_{\tilde{k}\tilde{p}} y_{qp}$:

$$y_{\tilde{k}\tilde{p}qp} \leq y_{qp} \qquad \forall q \in Q,\ \forall k, \tilde{k} \in K,\ \forall p, \tilde{p} \in P \tag{4.10}$$

$$y_{\tilde{k}\tilde{p}qp} \leq x_{\tilde{k}\tilde{p}} \qquad \forall q \in Q,\ \forall k, \tilde{k} \in K,\ \forall p, \tilde{p} \in P \tag{4.11}$$

$$y_{\tilde{k}\tilde{p}qp} \geq y_{qp} - (1 - x_{\tilde{k}\tilde{p}}) \qquad \forall q \in Q,\ \forall k, \tilde{k} \in K,\ \forall p, \tilde{p} \in P \tag{4.12}$$

Reformulating (4.9) using $y_{\tilde{k}\tilde{p}qp}$ yields:

$$\sum_{\tilde{k} \in K | d_{q\tilde{k}} \leq \bar{d}_q} \sum_{\tilde{p} \in P} e^{\left(u_{q\tilde{k}}\right)} y_{\tilde{k}\tilde{p}qp} = \sum_{k \in K | d_{qk} \leq \bar{d}_q} e^{\left(u_{qk}\right)} x_{kp} \qquad \forall q \in Q,\ \forall p \in P \tag{4.13}$$

4.2.2 Conceptual comparison

The approach to model customer choice of facilities differs significantly from previous models such as the Deviation flow refueling location model (DFRLM) proposed by Kim and Kuby (2011). In the DFRLM, all customers belonging to an OD-pair q are expected to take the shortest detour that gives them access to any active refueling facility. In the non-competitive setting discussed by Kim and Kuby (2011), this approach meets their requirements. However, for the desired competitive setting we do require a more granular approach. We introduce a distinction between facilities of different providers, and we no longer require a homogeneous facility choice from all customers traveling on the same OD-pair q. Instead, the proposed MNL model allows for heterogeneity in the decision making between customers of the same route.

Figure 4.1 highlights differences between the two approaches for a simplistic setting with a single OD-path q, and two deviations enabling the use of two (active) facilities (A,B) that are not directly located on q. The distance difference between a detour including A or B amounts to $d_{qA} = 0.5$ and $d_{qB} = 0.7$ respectively. In the DFRLM (4.1b), all customers willing to take a detour are assumed to homogeneously take the slightly shorter detour through A. As the additional distance to B is hardly noticeable, we argue that customers in practice might be somewhat indifferent between the two stations. We distinguish between the two providers (players) of the respective fuel stations, and divide the customer demand between both open facilities based on the relative utility they provide to customers traveling on q. If we assume equal attractiveness between facilities A and B ($a_A = a_B$), based on the shorter detour, A is awarded 62% of the total flow f_q and facility B expects a market share of 38%.

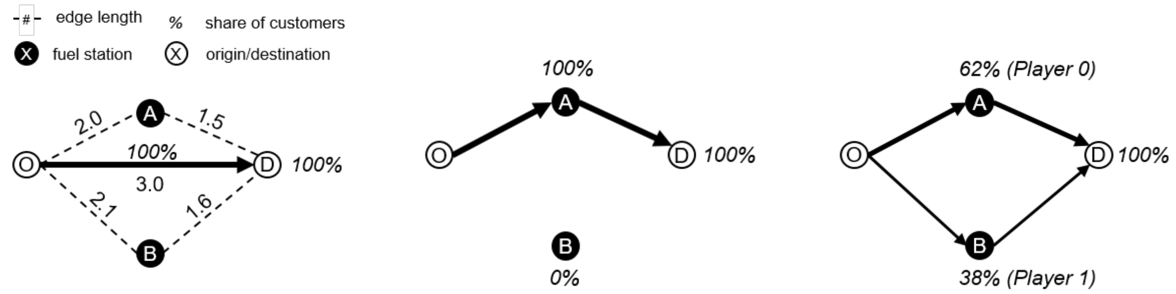

(a) Exemplary setup **(b)** DFRLM **(c)** C-FCLM

Figure 4.1: Exemplary comparison of different choice models ($\beta = 3, \bar{d}_q = 1.5$)

4.3 Solving the C-FCLM

To be able to solve the C-FCLM, we introduce the solution concept of Nash equilibria in combination with IPGs in Section 4.3.1. We discuss conditions for the existence of equilibria in Section 4.3.2, briefly summarize a solution method in Section 4.3.3 and define key performance indicators in Section 4.3.4.

4.3.1 Nash equilibria of integer programming games (IPGs)

The predominant solution concept for games modeling simultaneous competition between two or more players is the Nash equilibrium. A Nash equilibrium is a stable fix-point in which no player benefits from unilateral deviations (Nash, 1951). In a pure equilibrium, each player selects a distinct strategy, whereas in a mixed equilibrium players randomize between multiple different strategies. A strategic game in which player actions are defined by binary or integer decision variables can be represented as an IPG (Köppe et al., 2011). Based on this definition, the C-FCLM constitutes an IPG.

To represent the C-FCLM as an IPG, we rely on the following notation: We describe with X_p the non-empty and bounded action set of possible strategies of player p. X_p is given through the set of linear inequalities (4.2)-(4.3) as well as (4.8)-(4.13). Each player optimizes its own payoff Π_p, given through the objective function (4.1). We denote with φ a profile of strategies, where the tuple $\varphi = (\varphi_p, \varphi_{-p})$ consists of the player-specific strategy profiles. A player-specific strategy profile φ_p is comprised of elements $\varphi_{x_p} \in [0, 1]$, that denote the probability of player p playing strategy x_p. The expected payoff for player p in the strategy profile φ amounts to:

$$\Pi_p(\varphi) = \sum_{x_p \in X_p} \sum_{x_{-p} \in X_{-p}} \varphi_{x_p} \varphi_{x_{-p}} \Pi_p(x_p, x_{-p}) \tag{4.14}$$

A mixed Nash equilibrium is a profile of strategies $\sigma = (\sigma_p, \sigma_{-p})$ that satisfies:

$$\Pi_p(\sigma) \geq \Pi_p(x_p, \sigma_{-p}), \quad \forall p \in P,\ \forall x_p \in X_p \tag{4.15}$$

Note that the Nash equilibrium is a reasonable solution concept of the game, only if players have complete information of each others strategy space and payoffs. This assumption is reasonable for the C-FCLM: Due to space limitations in the examined inner-city setting and the required investments to build a fuel station from scratch it is rather unlikely that a fuel station provider would invest into a hydrogen fuel station outside of prior conventional fuel stations. As such, the strategy space of a provider is limited to their existing conventional fuel stations and any competitor could determine potential strategies outside-in. It is reasonable to assume that a fuel station provider could derive profit implications for each location selection of its competitor based on similar profit margins between the competitors and common knowledge of traffic flow patterns or customer count observations.

4.3.2 Properties of the C-FCLM

In the following, we determine the existence of Nash equilibria for the C-FCLM. As the decision variables of the C-FCLM (and therefore the action sets X_p for each player) are neither empty (the zero solution is a feasible solution) nor unbounded, it follows directly from Carvalho et al. (2022) that there exists at least one (mixed- or pure strategy) Nash equilibrium for any C-FCLM. To determine the existence of pure Nash equilibria, we draw on similarities to well-researched congestion games.

Theorem 4.1 (Existence of pure equilibria)**.** *A C-FCLM in which the allocation of customers to facilities per player is disjunct, the number of facilities per player is given,* *and the customers are indifferent between the facilities at their disposal, has a Nash equilibrium in pure strategies.*

Proof. The proof follows from the transformation of the C-FCLM into a congestion game. In a congestion game, P players compete for a set of resources $q = \{1, ..., Q\}$. Players choose a subset of all resources and face the cumulative cost of all selected resources. The cost c_q of an individual resource q is a non-decreasing function in the number of players n_q selecting the ressource q (Rosenthal, 1973).

Figure 4.2a visualizes how the choice of resources q (OD-trips) in the C-FCLM can be induced through the choice of facilities if we assume a disjunct allocation of customers to facilities (per player). By consequence, there is a one-to-one relationship between a strategy in the congestion game (selected resources/OD-trips) and a strategy in the C-FCLM (selected facilities that enable serving the respective OD-trips). If we further assume that customers are indifferent between the facilities that are able to serve them (i.e., $u_{qk} = u_{q\tilde{k}}$), and that the number of active facilities per player is fixed, we can define the cost function as:

$$c_q(n_q) = -m\frac{f_q}{n_q} \qquad \forall q \in Q \tag{4.16}$$

The total cost for player $p \in P$, serving customers $Q^p \subset Q$ amounts to:

$$c_p = \sum_{q \in Q^p} c_q = -m \sum_{q \in Q^p} \frac{f_q}{n_q} \underset{u_{qk}=u_{q\tilde{k}}}{\equiv} -m \sum_{q \in Q} f_q y_{qp} \qquad \forall p \in P \tag{4.17}$$

Under the assumption of a fixed number of active facilities per player, we can neglect fixed costs and thus this cost function is equivalent with the (negated) profit function (4.1). Thus, a C-FCLM with the aforementioned properties is a congestion game and by consequence has at least one pure equilibrium (Rosenthal, 1973). □

Note that this transformation into a standard congestion game is only possible under these restrictive assumptions. If customers are not indifferent between alternatives $u_{qk} \neq u_{q\tilde{k}}$, or the number of active facilities is not fixed, the transformation yields a weighted congestion game with separable preferences. For this class of congestion games, there is no guarantee for the existence of pure Nash equilibria (Milchtaich, 2009). In the general C-FCLM, where multiple facilities enable access to the same customers for one player (e.g., Figure 4.2b), a transformation into a congestion game is no longer possible as there

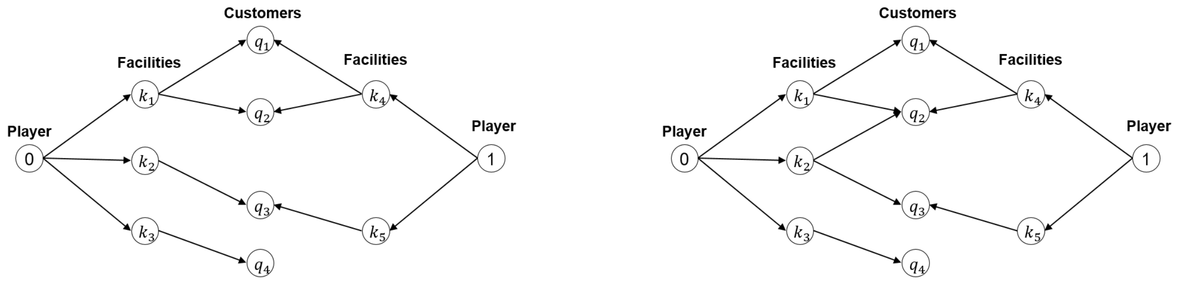

(a) C-FCLM with disjunct allocation of customers to facilities

(b) C-FCLM with overlapping allocation of customers to facilities

Figure 4.2: The C-FCLM as a congestion game: Choice of resources (customers q) is done indirectly through the choice of facilities.

is no longer a one-to-one relationship between choice of facilities and served customers (i.e., resources in the congestion game).

By consequence, we cannot guarantee the existence of pure-strategy Nash equilibria for the general C-FCLM. In contrast, one can imagine simple instances of the C-FCLM which do not permit a pure equilibrium (see Appendix 4.A). Note that although we cannot guarantee the existence of pure equilibria for the C-FCLM, in our practical application (Section 4.5), we always identified at least one pure NE.

4.3.3 Solving IPGs

In order to solve the C-FCLM we identify all (mixed an pure) equilibria of the IPG representation of the C-FCLM presented in Section 4.3.1. In contrast to a normal-form representation commonly used in game theory (Nisan et al., 2007), this IPG representation is beneficial as it does not require enumeration of all possible strategy combinations. This advantage is particularly useful for the C-FCLM, as already for small instances with 25 possible locations per player there are $(2^{25})^2$ potential strategy combinations. We use the eSGM-algorithm discussed in Chapter 3 to solve the IPG without full enumeration based on column-and-row generation. For convenience, we summarize key steps of the algorithm:

- Initialization: To initialize the algorithm, we limit each player to a very small subset of their original strategy set. For these limited strategy sets, we enumerate all strategy combinations and calculate the respective payoffs per player to derive a normal-form representation. This normal-form representation is referred to as the sampled game, as it only represents a small sample of the original IPG. For the

C-FCLM, we implement this initialization by restricting both players to a single strategy; the zero solution.

- Sample equilibria (master problem): We use an adaptation of the MIP-Nash algorithm (Sandholm et al., 2005) as the master problem. This master problem leverages a mixed-integer approach to identify all strategy combinations of the sampled game in normal-form, in which all players experience zero regret. By definition (Nash, 1951), these combinations represent the set of Nash equilibria of the sampled game.

- Best-responses (subproblem I): For each player and each equilibrium of the sampled game, best-responses are identified in a column-and-row-generating sub problem. In case there are best-responses, they are added to the master problem, enlarging the originally sampled game. This means the strategy space for players in the sampled game is extended by their respective best-responses. We repeat the identification of equilibria in the master problem (and subsequently repeat the identification of best-responses) after updating the normal-form representation of this extended sampled game.

- Additional equilibrium candidates (subproblem II): In case no further (unsampled) best-responses can be identified for the current equilibrium of the sampled game, this strategy combination is not only an equilibrium within the sample but also an equilibrium of the full IPG. However, it might well be that there are multiple equilibria, and the identified equilibrium is only one out of many. To guarantee the identification of all equilibria, the sampled game is further enlarged with additional equilibrium candidates in a second sub problem. An equilibrium candidate is defined as a strategy combination, in which no player benefits from deviating back to an already sampled strategy. This condition is a necessary, but not sufficient condition for any equilibrium of the IPG. Therefore, to determine whether the additional equilibrium candidates are indeed equilibria of the IPG, the identification of best-responses is repeated for the enlarged sampled game.

- Termination: The algorithm terminates under two conditions: There is no best-response to any identified equilibrium of the sampled game (i) and there are no additional equilibrium candidates (ii).

  Condition (i) guarantees, that the identified equilibria of the sampled game are equilibria for the full game (IPG): In case there is no best-response for any player to an identified equilibrium of the sampled game, no player benefits from unilateral deviation and the sampled game equilibrium constitutes a Nash equilibrium of the IPG.

Condition (ii) ensures, that all equilibria have been sampled and thus identified. In case there would be an equilibrium which is not part of the sampled game, it would constitute an equilibrium candidate and condition (ii) would not be fulfilled.

The described approach implies, that besides solving a mixed-integer master problem, we need to solve (4.1-4.13) repeatedly in the two subproblems to identify best-responses and candidate equilibria. Note that the MIP-Nash masterproblem (Crönert and Minner, 2021a; Sandholm et al., 2005) is linear for two players, quadratic for three players and features $n-1$ polynomials for n players. By consequence the algorithm does not scale well in the number of providers. We explore its scalability in number of candidate locations and the size of the transportation network in Section 4.4 and 4.5.

4.3.4 The price of equilibria and the value of the competitive solution

We can guarantee the existence of a mixed equilibrium for the C-FCLM (see Section 4.3.2). While this equilibrium is not necessarily unique, we can identify all equilibria using the described method (Section 4.3.3). Subsequently, an equilibrium selection approach (see for example Harsanyi (1995) or Harsanyi and Selten (1988)) could be used to single out a distinct equilibrium as the unique solution of the game. Instead, in our studies we opt to report aggregate solution results.

The Price of Anarchy (PoA) is commonly used in literature to measure the (in-)efficiency of equilibria in games. The PoA compares the objective value of the worst-case equilibrium with the optimal outcome if players would coordinate their efforts and share profits (Koutsoupias and Papadimitriou, 1999; Nisan et al., 2007). As such, it is a measure of inefficiency arising through selfish behavior and the failure to coordinate. Let us denote with Θ the set of all equilibria. Formally, we can define the PoA as:

$$PoA \stackrel{def}{=} \frac{\min_{\sigma \in \Theta} \sum_{p \in P} \Pi_p(\sigma_p, \sigma_{-p})}{\max_{x \in X} \sum_{p \in P} \Pi_p(x_p, x_{-p})} \tag{4.18}$$

with Π_p as the payoff for player p and σ as an equilibrium among the set of all equilibria Θ. The denominator denotes the maximum possible welfare as the cumulative payoff across all players, if players were to coordinate their actions $x = (x_p, x_{-p}) \in X$.

In contrast to the PoA, the Price of Stability (PoS) measures the best-case performance across all equilibria. In case the game only permits a single equilibrium, the PoS is equal to the PoA. In all other cases, the PoS is at least as close to 1 as the PoA (Nisan et al., 2007):

$$PoA \leq PoS \leq 1 \tag{4.19}$$

Formally, we define:

$$PoS \stackrel{def}{=} \frac{\max_{\sigma \in \Theta} \sum_{p \in P} \Pi_p(\sigma_p, \sigma_{-p})}{\max_{x \in X} \sum_{p \in P} \Pi_p(x_p, x_{-p})} \tag{4.20}$$

The PoS is of special interest when the involvement of third parties is considered. For example, a governmental policy could steer players towards the best-case equilibrium, and thus ensure a higher social welfare compared with the worst-case PoA. In the absence of such a third party involvement, the worst-case equilibrium can be just as likely as the best-case equilibrium. We therefore extend the PoS and the PoA with the Price of Equilibria (PoE) - a measure of equilibrium (in-)efficency in the mean case:

$$PoE \stackrel{def}{=} \frac{\sum_{\sigma \in \Theta} \sum_{p \in P} \Pi_p(\sigma_p, \sigma_{-p})}{|\Theta| \max_{x \in X} \sum_{p \in P} \Pi_p(x_p, x_{-p})} \tag{4.21}$$

The PoE is always in between PoA and PoS: $PoA \leq PoE \leq PoS \leq 1$. In the spirit of Roughgarden (2020), this allows us to avoid overly pessimistic (or optimistic) conclusions and to base our findings on the examination of average-case solutions.

There is no common measure, that is able to determine the benefit of addressing competition for a single decision maker. To highlight the benefits of a game theoretic model, that takes into consideration competitor actions, compared with a non-competitive model, we propose the Value of the competitive solution (VCS). The VCS follows the general concept of the (relative) Value of the stochastic solution (VSS) in stochastic programming. It compares the objective value between competitive and non-competitive model formulations for single decision makers in competitive settings. As such, it measures the losses a player has to face when ignoring apparent competition:

$$VCS_p \stackrel{def}{=} \begin{cases} 0 & , \Pi_p(\sigma_p, \sigma_{-p}) = 0 \\ 1 & , \Pi_p(\hat{x}_p, \sigma_{-p}) < 0 \\ \frac{\Pi_p(\sigma_p, \sigma_{-p}) - \Pi_p(\hat{x}_p, \sigma_{-p})}{\Pi_p(\sigma_p, \sigma_{-p})} & , \text{else} \end{cases} \qquad \forall p \in P \tag{4.22}$$

$$VCS \stackrel{def}{=} \frac{\sum_{p \in P} VCS_p}{|P|} \tag{4.23}$$

where $\Pi_p(\sigma_p, \sigma_{-p})$ is the payoff for player p, when p and the competitors $-p$ select their strategies according to the equilibrium σ. In contrast, $\Pi_p(\hat{x}_p, \sigma_{-p})$ is the payoff for player p when p plays the strategy $\hat{x}_p$ that is optimal in the absence of competition and the competitors $-p$ play according to the equilibrium σ. In general, the relative difference between the equilibrium solution and this non-competitive solution defines the VCS. Note that in our application, the equilibrium solution cannot be negative ($\Pi_p(\sigma_p, \sigma_{-p}) \geq 0$) as players are not forced to invest (the zero-solution is a feasible solution). To avoid division by zero, we separately deal with the case of a zero-solution in (4.22). However, it might very well be the case that the profit of the non-competitive solution is negative ($\Pi_p(\hat{x}_p, \sigma_{-p}) < 0$), if p did not foresee the action of competitors $-p$ and thereby made a non-profitable (and non reversible) investment. If the non-competitive solution is negative, the relative difference between the equilibrium solution and the non-competitive solution is bigger than 1. As the increase of this relative difference with a decreasing equilibrium solution ($\Pi_p(\sigma_p, \sigma_{-p})$) is counter-intuitive, we limit the VCS to 1. Equations (4.22)-(4.23) define the VCS for a single equilibrium. For our numerical experiments and the case study (Section 4.4 and 4.5) we report mean results across all equilibria.

4.4 Numerical results

The following numerical experiments are based on an implementation of the solution procedure introduced in Section 4.3.3 in Python 3. We use Gurobi 9.1 to solve the described master problem (identification of equilibria in the sampled game) and subproblems (extension of sampled game with best-responses and equilibrium candidates). We use a 2.6 GHz processor, with 4 available cores. Section 4.4.1 describes numerical experiments on a standard test instance commonly used for FCLMs, Section 4.4.2 determines the approximation error of the aforementioned clustering approach.

4.4.1 Test network

We test the C-FCLM on a 25-node network (Figure 4.3a), originally proposed by Simchi-Levi and Berman (1988) and used by various authors (Hodgson, 1990; Huang et al., 2015; Kim and Kuby, 2011; Kuby and Lim, 2005; Li et al., 2018; Lim and Kuby, 2010; Lin and

Lin, 2018; Tran and Nguyen, 2019) for flow-capture location models. Flow through the

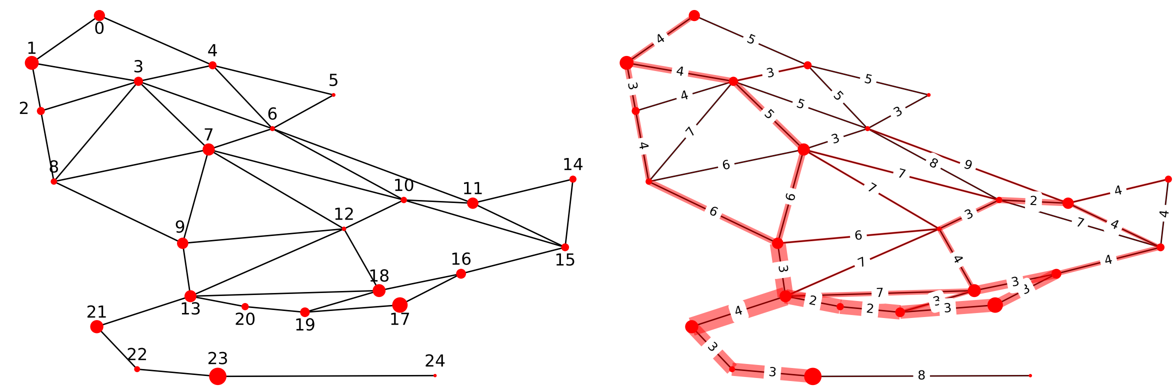

(a) Network nodes **(b)** Network edges and flow volume

Figure 4.3: Test network encompassing 25 nodes, 43 edges and 300 OD-pairs

network (see Figure 4.3b) is based on a gravitational model, assigning a flow quantity to each of the 300 OD-pairs on the shortest path between origin and destination based on their weights ξ (node size in Figure 4.3a, based on Hodgson, 1990):

$$f_q = \frac{\xi_i \xi_j}{d(i,j)^{1.5}} \qquad \forall q = (i,j) \in Q \tag{4.24}$$

We assume a willingness of customers to travel 10%, 20% or 30% longer than their original shortest path ($\hat{d} \in \{0.1, 0.2, 0.3\}$), without an absolute limit on the travel distance ($\bar{d} = \infty$). Apart from $\hat{d}$ (and by consequence $\bar{d}_q$), in our experimental setup we vary the investment cost per facility $c \in \{1500, 3000, 4500, 6000, 7500\}$. We assume $\beta \in \{1, 2, 3\}$ for the customer sensitivity towards driving distance. We neglect differences in facility attractiveness and assume $\alpha = 0$. Each player can choose from a set of $|K_p| \in \{2, 3, 4, 5\}$ randomly drawn locations out of the 25 nodes of the network. For each problem setup $(\hat{d}, c, |K_p|)$, we sample 100 candidate location sets. Table 4.1 reports mean results across all samples, highlighting effects on runtime, the number of equilibria $|\Theta|$, the VCS and PoA,PoE and PoS, depending on the examined parameter combination. As differences between different values of β were minuscule, results are aggregated for $\beta \in \{1, 2, 3\}$.

We observe that while the runtime stays strictly below 1 minute, it slightly increases with an increase in the willingness to travel of customers $\hat{d}$, and strongly depends on the number of candidate locations per player ($|K_p|$). This behavior is expected, as the increased strategy space of the players renders the proof that all equilibria have been identified more complex and time consuming. In contrast, the first equilibrium was identified in mere seconds. In all cases, recognizing competition improves the outcome for at least one player. This effect is particularly evident when customers are more sensitive with regards to deviation distance ($\hat{d} = 0.1$) and if investment costs are high (and by consequence, planning errors costly).

Figure 4.4 shows the effect of competition between red (circle) and green (triangle) for $c = 4500$, $\beta = 1$ and $\hat{d} = 0.1$ in detail. Selected locations are represented with filled markers, potential locations ($|K_p| = 4$) are shown with non-filled markers. If unaware of the existence of their competitor (Figure 4.4a), red chooses a location in the lower left, whereas green selects a location in the lower right. In contrast, if both players are aware of the emerging competition, green additionally opens a location in the lower left in direct competition with red. By acting in accordance with Hotelling's law (Hotelling, 1929), green can increase profits by 12% from 6247 (profit based on locations in Figure 4.4a) to 6991 (based on Figure 4.4b).

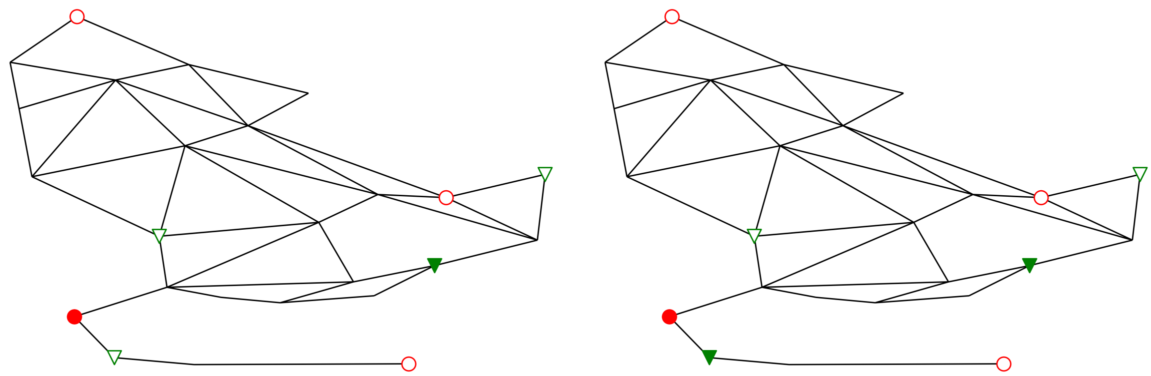

(a) Non-competitive setting **(b)** Competitive setting

Figure 4.4: Comparison of results in a competitive and a non-competitive setting for $c = 4500$, $\hat{d} = 0.1$, $\beta = 1$

For the examined 25-node network, the competitive solution of the C-FCLM leads to an average price of anarchy of 80% and a mean price of equilibria of 81% (across all parameter variations). This implies that in this case, central or coordinated planning can yield around 20% higher profits for fuel station providers. The majority of problem instances resulted in a unique Nash equilibrium, however some cases featured up to eight equilibria.

Table 4.1: C-FCLM results for the 25-node test network (mean of 300 samples per problem setup)

$\|K^p\|$	c	$\hat{d}$	runtime, sec	VCS, %	PoA, %	PoE, %	PoS, %	$\|\Theta\|$
2	1500	0.1	9.4	0%	92%	92%	92%	1
		0.2	10.7	0%	92%	92%	92%	1
		0.3	12.4	1%	91%	91%	91%	1
	3000	0.1	11.2	4%	94%	94%	94%	1
		0.2	18	4%	87%	87%	87%	1
		0.3	19.5	3%	82%	82%	82%	1
	4500	0.1	12.3	7%	88%	90%	92%	1.2
		0.2	14.7	5%	91%	92%	93%	1.2
		0.3	17.8	1%	91%	91%	91%	1
	6000	0.1	10.8	8%	74%	77%	81%	1.2
		0.2	14.7	4%	84%	86%	87%	1.2
		0.3	15.6	1%	86%	87%	89%	1.2
	7500	0.1	10.1	38%	86%	91%	96%	1.2
		0.2	14.1	15%	72%	74%	76%	1.2
		0.3	14.5	2%	79%	79%	79%	1
3	1500	0.1	13.9	1%	92%	92%	92%	1
		0.2	21.4	2%	89%	89%	89%	1
		0.3	26.4	4%	85%	85%	85%	1
	3000	0.1	15.7	3%	86%	87%	87%	1.2
		0.2	21.4	3%	88%	88%	88%	1
		0.3	31	3%	80%	80%	80%	1
	4500	0.1	14.7	4%	75%	77%	79%	1.4
		0.2	21	2%	80%	81%	83%	1.4
		0.3	34.2	3%	78%	81%	84%	1.6
	6000	0.1	13.3	12%	66%	68%	71%	1.2
		0.2	18.6	4%	76%	76%	76%	1
		0.3	27.2	4%	78%	78%	78%	1
	7500	0.1	12.3	28%	71%	78%	81%	1.2
		0.2	17	7%	65%	65%	65%	1
		0.3	20.1	2%	70%	70%	70%	1
4	1500	0.1	21.4	0%	93%	93%	93%	1
		0.2	34.2	6%	87%	87%	87%	1
		0.3	44	8%	83%	83%	83%	1
	3000	0.1	23.6	3%	88%	89%	90%	1.4
		0.2	41.7	5%	80%	81%	82%	1.4
		0.3	41.4	3%	79%	79%	79%	1
	4500	0.1	26.1	5%	81%	83%	85%	1.4
		0.2	26.7	4%	80%	80%	80%	1
		0.3	33.4	3%	81%	81%	81%	1
	6000	0.1	17.7	6%	75%	75%	75%	1
		0.2	31	5%	78%	78%	79%	1.4
		0.3	34.6	7%	77%	77%	77%	1.4
	7500	0.1	16.4	21%	74%	74%	74%	1
		0.2	22.3	6%	68%	68%	68%	1
		0.3	29.2	4%	71%	72%	73%	1.2
5	1500	0.1	34.3	1%	90%	90%	90%	1
		0.2	58.9	4%	84%	84%	84%	1
		0.3	65.2	8%	81%	81%	81%	1
	3000	0.1	28.6	3%	85%	85%	85%	1
		0.2	45.4	3%	81%	81%	82%	1.2
		0.3	55	5%	77%	77%	77%	1
	4500	0.1	32.6	2%	77%	77%	78%	1.2
		0.2	40.9	1%	79%	81%	81%	1.4
		0.3	56.3	6%	74%	76%	77%	1.4
	6000	0.1	27.8	3%	72%	72%	72%	1
		0.2	37.4	2%	75%	75%	75%	1.2
		0.3	49.8	3%	70%	71%	72%	1.4
	7500	0.1	24.6	13%	62%	62%	62%	1
		0.2	32	0%	70%	70%	70%	1
		0.3	37.2	2%	70%	70%	70%	1
		mean	26.4	5%	80%	81%	82%	1.1

4.4.2 Clustering

In Section 4.4.1 we observed influences on runtime through the number of candidate locations $|K_p|$ and the relative deviation distance $\hat{d}$. In addition to these previously examined runtime effects, the number of distinctly modeled OD-trips $|Q|$ clearly drives the number of variables in the C-FCLM and by consequence model complexity. While the examined 25-node network can be solved (i.e., all equilibria are identified) within seconds for most cases, real-world applications might lead to significantly more complex networks with a large number of OD-pairs $|Q|$ (see Section 4.5) that will exceed computational limitations during the identification of all equilibria.

As the distance from one node to a neighboring node can be negligibly small in such a dense urban environment, we expect only slight approximation errors when aggregating similar OD-trips. We propose to combine OD trips, that hardly differ (in terms of euclidean distance) in their origin and destination node, into a common cluster. This cluster is represented through a single representative OD pair (the cluster center), which accounts for the cumulative flow of all OD-pairs within the cluster. We aggregate all OD-pairs $q \in Q$ into $|\mathcal{C}|$ clusters $c \in \mathcal{C}$ minimizing the total weighted squared error Φ:

$$\Phi = \sum_{\acute{q} \in Q} \min_{c \in \mathcal{C}} f_{\acute{q}} ||\acute{q} - c||^2, \tag{4.25}$$

with $\acute{q} \in \mathbb{R}^4$ as a four dimensional vector representation of the OD-pair q with x, y-coordinates for the origin and the destination node. Finding the global optimum to (4.25) is NP-hard even for non-weighted clustering in $\mathbb{R}^2$ (Mahajan et al., 2012) or when limiting the number of clusters $|\mathcal{C}|$ to two (Aloise et al., 2009). Therefore local search algorithms such as kmeans (Lloyd, 1982) are used. We use kmeans++ (Arthur and Vassilvitskii, 2006), an extension of the original kmeans algorithm that ensures $\mathcal{O}(\log |\mathcal{C}|)$-competitiveness with optimal clustering. In the proposed clustering approach we use the coordinates of the origin and the destination node, as well as the flow f_q to fully characterize a given OD-pair q and to determine its proximity to other OD-pairs. Naturally, one could also think of alternative proximity measurements to determine which OD-trip belongs to which cluster. For example, we could define clusters by enforcing that each OD-trip in the cluster shares a certain percentage of its shortest path with the shortest path of the cluster center. However, this definition would imply that two OD-trips that start and end very close to each other, and whose shortest paths are close but separate throughout the trip would not be part of the same cluster, despite facing similar deviations to potential fuel station locations. Another alternative proximity formulation could be based directly on these deviation distances to potential fuel station locations (d_{qk}). Two OD-trips, that face the same or similar distances to all potential fuel station locations $k \in K$ can clearly be considered similar and should belong to the same cluster. While this would be an appropriate proximity measure in theory, it would imply clustering in $\mathbb{R}^{|K|}$ instead of $\mathbb{R}^4$ and thus significantly hinder computational performance. We therefore opted to use origin/destination coordinates to determine proximity of OD-trips.

We apply this clustering approach to examine the introduced error in a controlled numerical experiment based on the 25-node test network. Clearly, this error depends on the allowed number of clusters. The optimal cluster count $|\mathcal{C}|$ is a trade-off between reduced computation time and the resulting approximation error. To evaluate this trade-off, we determine the difference between results of the original network ($|Q| = 300$) and the clustered representation while varying $|\mathcal{C}|$. As in Section 4.4.1, player specific candidate locations K_p are drawn at random from a uniform distribution across all 25 nodes. To isolate clustering effects, we keep c and $\hat{d}$ fixed ($c = 6000$, $\hat{d} = 0.1$).

Figure 4.5 shows the effects of clustering on the absolute error Δ_{abs} for PoE and VCS. Δ_{abs} is the absolute difference (in PoE or VCS, respectively) between the original solution and the solution of the clustered representation (evaluated in the objective functions of the original model). Each data point reflects the mean of 100 samples. When cutting the number of explicitly modeled OD-trips in half ($|\mathcal{C}| = 150$), we can reduce runtimes by more than 50%, at the cost of deviations of 2-4 percentage points in the reported VCS and PoE. Note that in this synthetic setting, the ratio of candidate locations across both players ($|K| = 6$ to 8) to the total number of nodes in the network (25) is high. In reality, if players are restricted to a small subset of locations (e.g., existing conventional fuel stations) in a very large and dense urban network, we expect smaller deviations.

4.5 Case study

In the following case study, we assess HFS location selection for two competing players in the city of Munich, Germany. Reported runtimes have been achieved with a 2.6 GHz processor at 16 available cores.

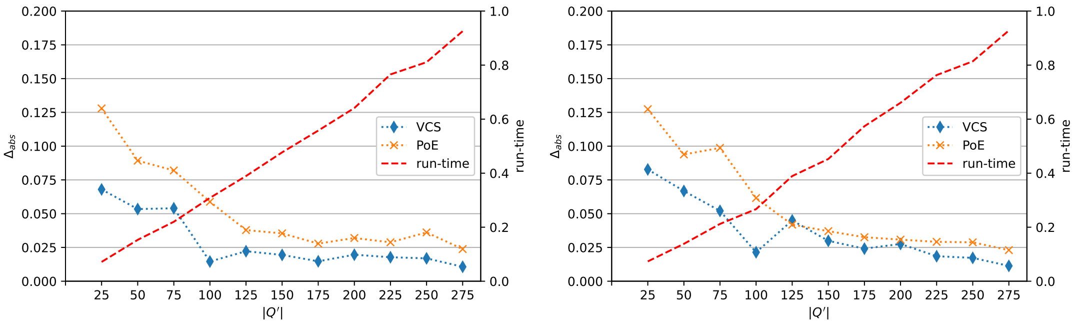

(a) $|K_p| = 3$ **(b)** $|K_p| = 4$

Figure 4.5: Effects of clustering OD-trips on PoE, VCS and runtime (normalized for $|Q|$=300), mean results for 100 samples per data point

4.5.1 Pre-processing

Figure 4.6 shows the network of main roads in Munich. This network consists of more than 9000 nodes and more than 15000 edges (Figure 4.6a). Each edge refers to a road section with a mean length of 146 meters. On average, a node connects 2.4 road sections. Flow through the network in OD-pairs is given by agent-based simulations from Moeckel et al. (2020) (Figure 4.6b).

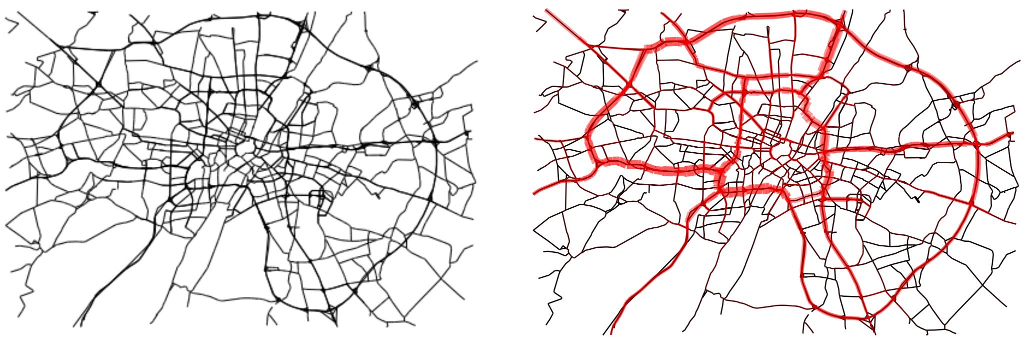

(a) Munich main road network (>9000 nodes, >15000 edges) **(b)** Flow density through the network (>50000 OD-pairs)

Figure 4.6: Munich main road network and traffic flow

The pre-processing builds on the following assumptions:

(1) As nodes representing junctions of the road network are always superior (in terms of vehicle flow) to neighboring nodes that are not part of a junction and the distance to the next junction is negligible within a city, we can neglect non-junction nodes.

(2) As the real estate in Munich is limited, and investment costs for building a new fuel station are prohibitively high, we expect hydrogen fuel stations exclusively at existing (conventional) fuel stations of the respective players. To reduce the set of candidate locations to existing fuel stations, we rely on publicly available geo-spatial information from openstreetmap (OSM, 2020). We parse all known locations of existing fuel stations and for each select the closest junction as a proxy location. We limit our examination to the two biggest fuel station providers in Munich (P_0=Shell, P_1=Aral). Figure 4.7a shows the candidate locations for both players.

(3) As many OD-trips through the network share similar paths on the majority of their route, we can reduce model complexity through aggregation. We depict multiple similar OD-trips with a single (representative) OD-pair q (see Section 4.4.2) and reduce the number of explicitly modeled OD-pairs into $n_c = 1000$ clusters. Figure 4.7b shows an exemplary representation of a cluster through its center route (red) with the original OD-trips in blue. The final model consists of $|Q| = 1000$ OD-pairs and $|K| = 34$ potential facility locations among $|N| = 2$ competing players.

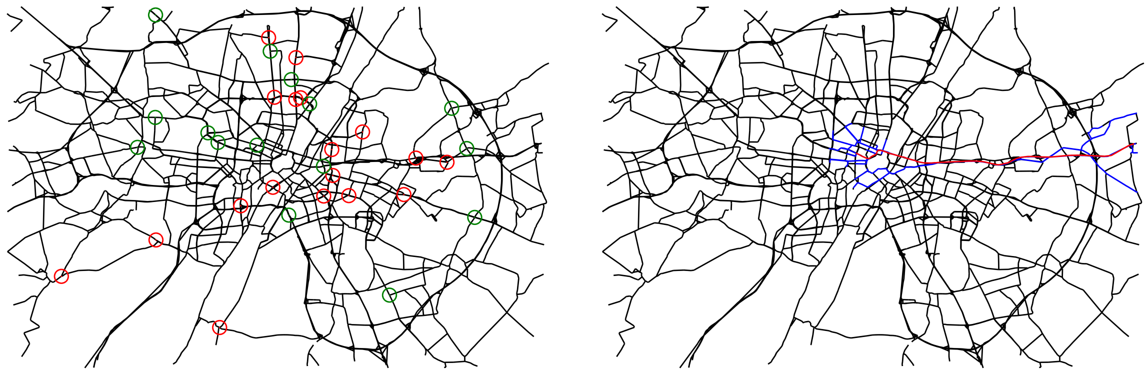

(a) Potential fuel station locations for P_0 (red) and P_1 (green)

(b) Exemplary cluster of original OD-trips (blue) and center representation (red)

Figure 4.7: Potential fuel station locations and exemplary cluster representation of OD-trips

We aggregate the traffic data to one representative year and accumulate all (potential) customers sharing the same origin and destination to the flow f_q of said OD-pair q. We denote with γ the share of all cars using an FCEV-drivetrain (market penetration) and scale the traffic flow with the expected FCEV market penetration according to three different scenarios ($\gamma \in \{5\%, 10\%, 15\%\}$), as well as the average hydrogen consumption per FCEV to arrive at actual hydrogen fuel demand data. We neglect differences in size, price or brand between (potential) fuel stations ($\alpha = 0$). Based on Brown et al. (2013), we assume a retail markup of $\approx$ 10 ct per kg (above the current markup for conventional fuel) and 20 ct per kg of additional margin through convenience product sales in store. At an average transaction volume of $\approx$ 3.3 kg, this yields $m = 1$ EUR in operational margin per served customer.

The assumption that differences in prices and profit margins between providers are negligible is common in prior literature on competitive hydrogen fuel station location selection (Bersani et al., 2009) or in prior competitive variants of the FCLM (Berman and Krass, 1998; Wu and Lin, 2003). This assumption is sensible, as the majority of the contribution margin of fuel station providers comes from the sales of convenience goods rather than fuel sales (Brown et al., 2013). We expect the margin of convenience goods to be comparable across fuel stations and brands. Neglecting potential economies of scale, the margin per customer m can thus be assumed to be exogenous industry knowledge. Nevertheless, price competition clearly plays a role in the day-to-day business for conventional fuel stations (Bergantino et al., 2020) – and thus might also be relevant for hydrogen fuel stations in the near term. However, while the number of low-cost competitor stations in close proximity has a measurable effect on price levels for conventional fuel stations, other, structural factors such as type of road, population count, real estate value or the number of commercial businesses are by fare more influential on the fuel price (Bergantino et al., 2020). Hence price competition will likely be limited to operational decision making and will not influence the strategic location decision on the HFS network as a whole.

Depending on the fuel station capacity and technology, Apostolou and Xydis (2019) list investment costs for HFS ranging from $1.0 - 2.4$mn EUR, for our calculations we assume an initial investment of 1.5mn EUR and depreciate over 15 years with a yearly depreciation of 0.1mn EUR and maintenance costs of 0.05mn EUR yielding yearly fix costs of $c = 0.15$mn EUR.

As we are lacking data on actual customer sensitivities (e.g., $\beta, \hat{d}$), we explore results under various problem settings using a full factorial design. Besides the scenarios of FCEV market penetration (γ), we vary the maximum deviation distance customers are willing to drive ($\hat{d} \in \{10\%, 15\%, 20\%\}$) and the customer sensitivity towards distance $\beta \in \{1, 2, 3\}$. In addition to the relative deviation distance limit $\hat{d}$ we employ an absolute limit of $\bar{d} = 3$ as we do not expect drivers taking an inner-city detour of more than 3 km to refuel. In total, this setup yields 27 problem instances with different parameter combinations $(\beta, \hat{d}, \gamma)$.

4.5.2 Computational results

In all 27 instances, we identified at least one pure Nash equilibrium, even though the existence of pure Nash equilibria is not guaranteed for the C-FCLM (see Section 4.3.2). As practical implications are hard to deduct from mixed equilibria, we report results for pure equilibria only. In some instances, the equilibrium was not unique, with up to 3 equilibria per instance. In these cases with multiple equilibria, we report mean values across all (pure) equilibria. Figure 4.8 shows the effects of the different factor levels on the average yearly profit per provider ($\tilde{\Pi}_p$), the VCS (see Section 4.3.4), the total number of active hydrogen fuel stations n and the total solution runtime.

We observe that solution runtime increases with an increase in $\hat{d}$, as customers take more alternative stations into account. Solution runtime also increases with a higher penetration: As it becomes more attractive to open facilities for fuel station providers, the average number of hydrogen stations n increases and the strategy space expands. In the most complex setting with a high willingness to travel $\hat{d} = 0.2$ and a large hydrogen penetration $\gamma = 0.15$ runtime reaches up to 15000 seconds ($\approx 4.2h$), despite prior clustering of OD-trips. Evidently, a higher penetration increases the average annual profit per provider from $\approx$ 13k EUR at $\gamma = 5\%$ to more than 200k EUR at $\gamma = 15\%$. In contrast, the competitiveness (i.e., the VCS) is especially high for a very small FCEV penetration ($\gamma = 5\%$). As margins are especially small with low FCEV adoption, two competing HFS that are unintentionally in very close proximity might very well yield negative returns for both providers, highlighting the need to take possible competitor actions into consideration during location planning. On average, across all instances the VCS amounts to 17% which suggests that hydrogen fuel station providers who take into account likely competitor (re-)actions can substantially increase profitability.

If the two players were allowed to collaborate, they would be able to further increase profits by 28% (based on a mean PoE of 72%). Notably, this increase in provider profit comes at the cost of customer utility: Customers would suffer from an increase in deviation distance of up to 10% (average: 3%) as collaboration between providers would lead to a reduced number of HFS. For policymakers this means, that while fostering collaboration (e.g., through less strict antitrust laws or government-backed provider associations) benefits the providers, it does not imply a higher number of fuel stations and thus could have an adverse effect on the customer (i.e., FCEV driver) experience.

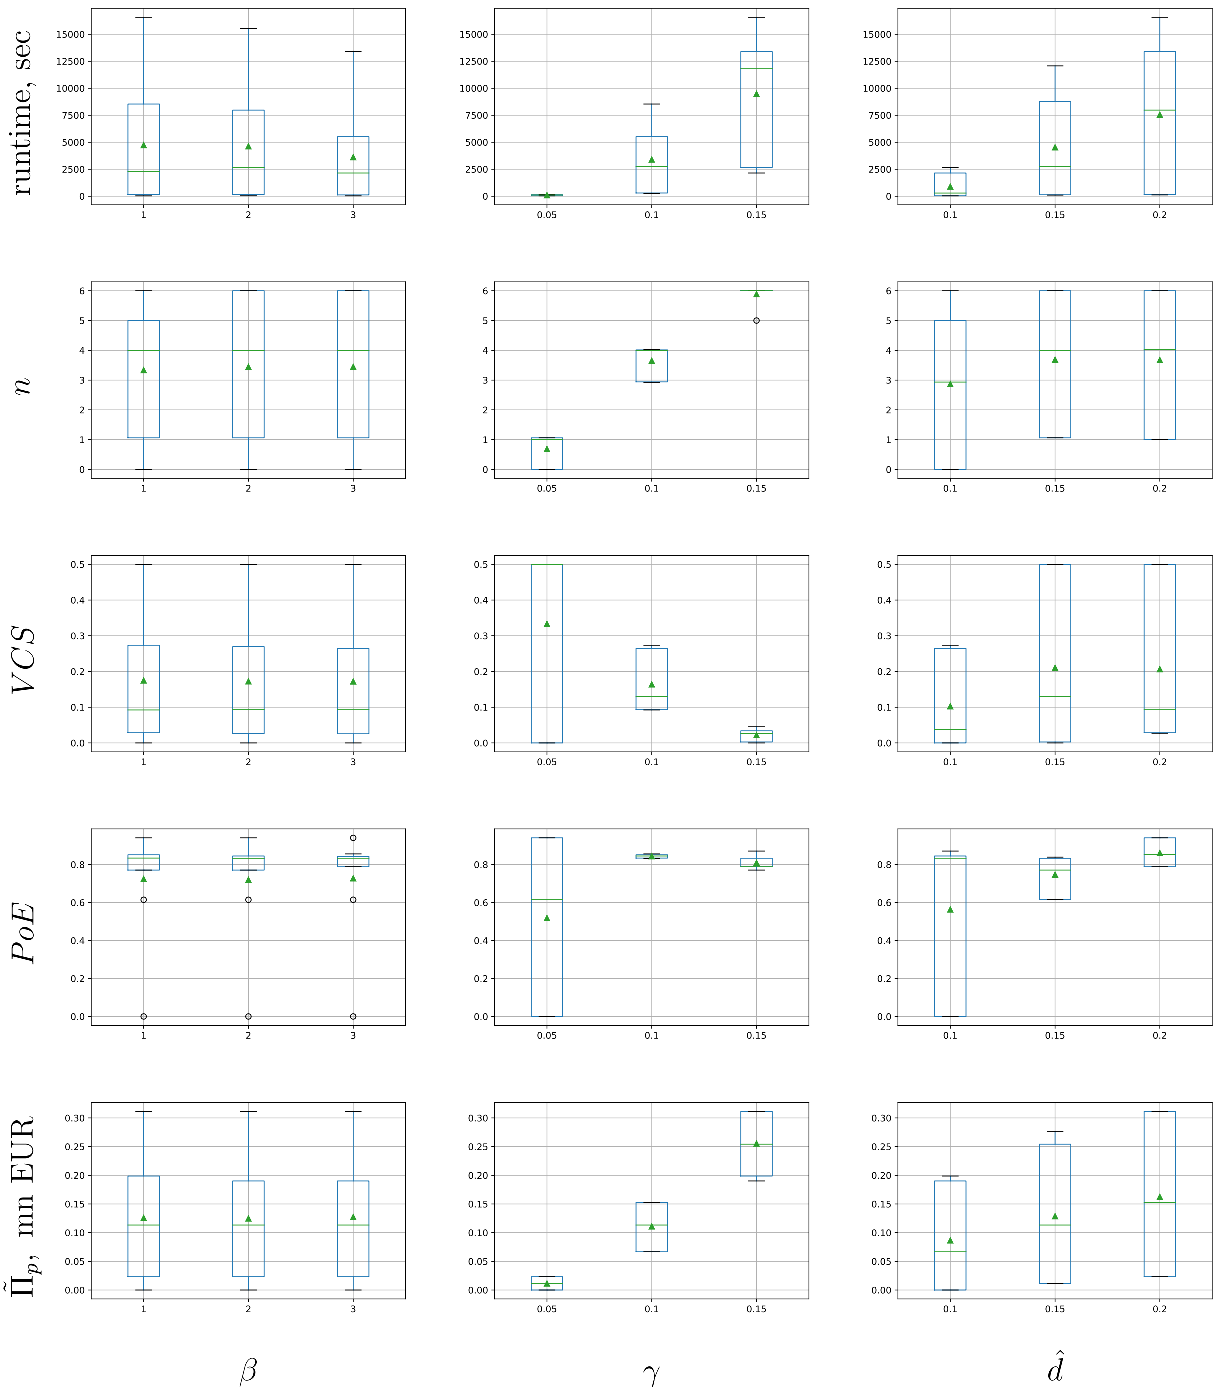

Figure 4.8: Effects on runtime, average number of fuel stations n, the VCS and PoE, and the average (yearly) payoff per player $\tilde{\Pi}_p$ for different parameter settings of β, γ, $\hat{d}$

Figure 4.9 exemplifies spatial differences between a competitive and a non-competitive setting for $\gamma = 15\%$, $\hat{d} = 10\%$ and $\beta = 3$. In the non-competitive setting (Figure 4.9a), where both players are unaware of competition and optimize independently, we see an accumulation of two competing HFS in the east. This close proximity significantly reduces profitability for the red player, leading to losses when compared to the spatially differentiated competitive solution in Figure 4.9b.

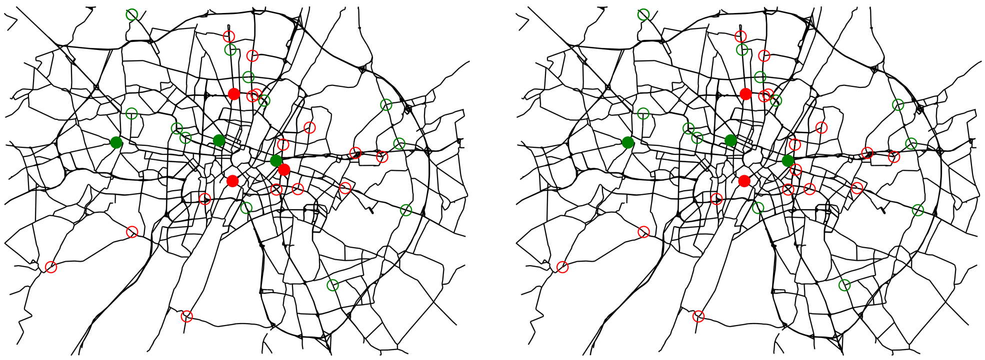

(a) Non-competitive setting

(b) Competitive setting

Figure 4.9: Comparison of results in a non-competitive and in a competitive setting for $\gamma = 15\%, \hat{d} = 10\%, \beta = 3$

4.6 Conclusions and future research

By extending the FCLM to a competitive setting, and solving the resulting model in various numerical experiments and in a case study, we show that taking emerging competition and possible competitor actions into account can help decision makers who select hydrogen fuel station locations to realize savings of 17% on average (mean across all examined scenarios). A price of anarchy around 60% and a mean price of equilibria of 72% suggests, that collaboration between providers (under applicable antitrust laws) or a centrally planned hydrogen supply chain could substantially benefit providers. While this might decrease the hurdle to invest at very low FCEV market penetration, it does come at the cost of longer deviation distances for customers at higher market penetration, with customers experiencing a sparser network of HFSs.

The proposed model assumes equal prices and margins for all competitors. Evidently, a player could try to compensate an inferior location by partially decreasing fuel prices at the cost of margins and thus increase the facility attractiveness a_k. A natural extension of the model would hence reflect not only decisions on locations but also the price level of each location and incorporate the effect of different price levels among facilities on customer choice. Further research opportunities lie within a multistage stochastic extension of the proposed framework for competitive flow-capture location-allocation models. As we are facing uncertain technology adoption, the advantages of a dynamic, stochastic formulation are evident. For non-competitive settings, similar formulations are readily available (Hosseini and MirHassani, 2015; Tan and Lin, 2014; Wu and Sioshansi, 2017). In a competitive setting, however, such an extension is not straightforward as characteristics of dynamic and stochastic games (e.g., subgame perfection and non-credible threats, see Selten, 1975) have not yet been properly addressed for IPGs and conventional (extensive form-based) game theoretic models able to identify subgame perfect equilibria would require the full enumeration of each and every possible strategy combination.

Appendix 4.A C-FCLM without a pure equilibrium

Imagine a C-FCLM with four OD-trips $q \in Q = \{q_1, q_2, q_3, q_4\}$ and $|K| = 4$ facilities split on $|P| = 2$ competitors: $K_0 = \{k_1, k_2\},\ K_1 = \{k_3, k_4\}$. The utility u_{qk} of a facility k perceived by customers on a certain OD-trip q is given in Table 4.2. Note that the utility $u_{q_4k_1}$ equals zero, we assume that q_4 cannot be served by k_1 as $d_{q_4k_1} > \bar{d}_q$. We further assume customer flows of $f_q = [215, 220, 1, 2]$, margins and fix costs of $m = 1$ and $c = 100$, respectively. Given the small number of facilities per player, we can enumerate all strategies, and transform the C-FCLM into the normal-form game of Table 4.3. Deletion of dominated strategies yields Table 4.4, which shows no pure equilibria and one mixed equilibrium with player one randomizing between k_1 (70%) and k_2 (30%) and player two randomizing between k_3 and k_4 with 73% and 27% probability.

Table 4.2: Utility u_{qk} for examplary C-FCLM without pure Nash equilibrium

u_{qk}	k_1	k_2	k_3	k_4
q_1	1.1	3.2	1.4	3.6
q_2	2.8	0.7	2.7	0.7
q_3	0.7	1.1	1.1	0.7
q_4	0	1.8	1.6	0.7

Table 4.3: C-FCLM without pure Nash equilibrium in normal-form representation

		$x_1 = [x_{k3}, x_{k4}]$			
		$[0, 0]$	$[1, 0]$	$[0, 1]$	$[1, 1]$
$x_0 = [x_{k1}, x_{k2}]$	$[0, 0]$	0,0	0,338	0,338	0,238
	$[1, 0]$	336,0	107.3,130.8	113.5,124.5	23.1,114.9
	$[0, 1]$	336.0,0	112.9,125.1	98.1,139.91	7.2,130.8
	$[1, 1]$	238,0	110.0,28.2	92.7,45.3	1.8,36.2

Table 4.4: C-FCLM without pure Nash equilibrium in reduced normal-form representation

		$x_1 = [x_{k3}, x_{k4}]$	
		$[1, 0]$	$[0, 1]$
$x_0 = [x_{k1}, x_{k2}]$	$[1, 0]$	107.3,130.8	113.5,124.5
	$[0, 1]$	112.9,125.1	98.1,139.9

Chapter 5

Inverse optimization for parameter estimation arising from competitive retail location selection

When determining store locations, competing retailers must take customers' store choice into consideration. Customers predominantly select which store to visit based on price, accessibility, and convenience. Incumbent retailers can estimate the weight of these factors (customer attraction parameters) using granular historical data. Their location decision under full information and simultaneous competition translates into an integer programming game. Unlike those incumbents, new entrants lack this detailed information; however, they can observe the resulting location structure of incumbents. Assuming the observed location structure is (near-)optimal for all incumbent retailers, a new entrant can use these observations to estimate customer attraction parameters. To facilitate this estimation, we propose an inverse optimization approach for Integer programming games (IPGs), enabling a new entrant to identify parameters that lead to the observed equilibrium solutions. We solve this inverse IPG via decomposition by solving a master problem and a subproblem. The master problem identifies parameter combinations for which the observations represent (approximate) Nash equilibria when compared with optimal solutions enumerated in the subproblem. This row-generation approach extends prior methods for inverse integer optimization to competitive settings with (approximate) equilibria.

We compare the decision-making of new entrants selecting locations based on expected values or scenarios of customer attraction parameters with new entrants using inversely estimated parameters for their location decisions. New entrants who rely on inversely optimized parameters can improve their profits by 4-11% on average. This benefit can be realized with as little as one or two observations, yet additional observations help to increase prediction reliability significantly.

5.1 Introduction

By offering lower prices, some grocery retailers established large stores (superstores) in the outskirts of cities to attract consumers to shop at these stores by purchasing larger quantities with less frequent visits (Bell et al., 1998; Leszczyc et al., 2004). Nowadays, increasingly fast-paced consumers find more convenient alternatives, moving to bite-size consumption in smaller stores close to their work or home (Belavina et al., 2017). Retailers react accordingly and prioritize stores inside cities, focusing on customer proximity but reducing sales area and assortment to compensate for the higher cost of retail space.

While store location is a key element of customer attraction, other factors such as store brand recognition, the retailers' assortment, or price levels can also play a crucial role in the customers' store selection. Without in-depth knowledge of these customer attraction parameters, retailers cannot choose their store locations optimally. Prior research determines these customer attraction parameters based on customer surveys which are prone to response biases, or using estimation techniques based on discrete choice models in combination with granular sales data (e.g., maximum likelihood estimation). Ben-Akiva and Lerman (1997) provide a general introduction to discrete choice theory, for applications to retail operations and market shares, see Cooper and Nakanishi (1988) or Berbeglia et al. (2021). While maximum likelihood estimation and similar approaches can be an appropriate tool for incumbent retailers with plenty of sales records available at existing locations; a new entrant is unlikely to have access to such granular data. To overcome this challenge, we address this issue faced by a new entrant by examining competitive retail location selection and estimating customer attraction parameters based on limited, observable information on existing locations of incumbents.

Selecting store locations of multiple competing retail chains is modeled as a simultaneous competitive location problem. This assumption on the sequence of decision-making between the competing retail chains is a key differentiator between competitive location models (Eiselt et al., 2019). In line with location selection of established retail chains in practice, we focus on store location selections arising from a simultaneous competitive situation between a finite number of competing retail chains, with no clear leader-follower relationship. We model such a simultaneous competitive setting as an IPG. An IPG is a non-cooperative, full-information simultaneous move game between two or more competing players in which the (integer) decision of each player affects the objective function of their competitors. The predominant solution concept for IPGs is a Nash equilibrium (Carvalho et al., 2022; Köppe et al., 2011; Nash, 1951); for our application, such an equilibrium is a combination of store location decisions in which no retail chain benefits from unilateral deviation.

We consider the situation of a single outside party who can observe the result of a non-cooperative simultaneous location selection game (formulated as an IPG) associated with multiple incumbents. In the following, we assume this external observer is a new entrant who is interested in extracting information to optimize its own market entry location selection. Similarly, public authorities and city planners as external observers could be interested in the parameter estimates to guide regulatory and planning efforts.

Incumbents have full knowledge of each other's payoffs and the customer choice behavior (e.g., through estimates based on historical sales data). They know the customers' valuation for brand, store location, and travel distance or convenience of general accessibility. Their observable location decisions are thus (near-)optimal and constitute an (approximate) Nash equilibrium of an IPG. However, in contrast to the incumbents, the new entrant does not know the customers' choice behavior and lacks the sales data required for conventional estimation. Yet, they can use the observed (near-)optimal location selection of incumbents in combination with their knowledge of other parameters (e.g., location costs, population count) to deduce information on customer store choice based on inverse optimization.

Given a (feasible) solution to the IPG associated with the incumbents (i.e., the forward problem), the new entrant solves an inverse optimization problem (inverse IPG). The objective of this inverse IPG is to find a set of parameter values such that a variation of the forward problem IPG based on these parameters leads to optimality of the provided solution. Therefore, our goal is to examine the new entrant's problem by solving the inverse optimization of IPGs; i.e., the estimation of the parameter set that defines the observed equilibrium arising from competition between the incumbents. Subsequently, these estimated customer attraction parameters improve the ability of the new entrant to make their own market entry location decisions.

5.1.1 Contribution

While existing approaches can inversely optimize parameters in integer programs (Wang, 2009) or in continuous equilibrium applications (Bertsimas et al., 2015), no approach that integrates both integrality and competition (equilibria) in inverse optimization exists.

We contribute to the literature on inverse optimization by extending prior methods for inverse integer problems to solve the inverse optimization of competitive integer problems (IPGs). This enables the estimation of parameters that explain the behavior of multiple competitors observed in a Nash equilibrium. To deal with noisy observations, which do not fully represent the decision-making expected based on the underlying ground-truth parameters, or with situations in which no pure Nash equilibrium exists, we introduce the method for the more general $\epsilon-$Nash equilibria. The proposed method offers a plethora of applications beyond competitive location decisions, including the estimation of parameters in capacity or inventory games, as well as (competitive) investment or assortment decisions. We thus provide an approach to evaluate (integer) decisions under simultaneous competition observed in practice.

5.1.2 Organization

Section 5.2 provides background on IPGs and formalizes the (inverse) problem. We examine a stylized example and develop a general solution methodology in Section 5.3. In Section 5.4, we conduct numerical experiments to assess the value of the developed methodology for new entrants and its ability to extract information from observations. Section 5.5 summarizes the main findings and provides an outlook on future research directions.

5.2 Context and problem setting

We briefly describe IPGs in Section 5.2.1 and introduce the general problem and key assumptions in Section 5.2.2, before formalizing the customer choice problem (Section 5.2.3), the incumbents location selection problem (Section 5.2.4) and the inverse optimization problem of the new entrant (Section 5.2.5).

5.2.1 A general description of integer programming games (IPGs)

An IPG is a simultaneous move game between n-players $i \in I = \{1, ..., n\}$ in which a player's strategy $\mathbf{x}_i \in S_i$ is comprised of z^i bounded integer decisions (Köppe et al., 2011). The finite action set S_i of possible strategies for each player i is given by Q^i inequalities:

$$S_i = \left\{\mathbf{x}_i \in \mathbb{Z}^{z_i} \mid w_q(\mathbf{x}_i) \leq 0,\ \forall q \in \{1, \ldots, Q^i\}\right\}. \tag{5.1}$$

Here, w_q represents any real-valued function. Prior literature (Carvalho et al., 2022; Köppe et al., 2011) assumes w_q to be linear, but, in the case of Carvalho et al. (2022), does not restrict all decision variables to integers. We denote with $-i$ all players except i, i.e., $-i = I \setminus \{i\}$. Then, $\mathbf{x}_{-i}$ is a combination of strategies of i's competitors, i.e., $\mathbf{x}_{-i} = (\mathbf{x}_{\tilde{i}})_{\tilde{i} \in -i}$. Players i choose their strategies $\mathbf{x}_i$ to maximize payoffs $\Pi_i(\mathbf{x}_i, \mathbf{x}_{-i})$, which depends both on their own decision and the decision of their competitors. Thus, a popular solution concept for IPGs is a Nash equilibrium in which no player benefits from unilateral deviation (see Nash, 1951):

$$\Pi_i(\mathbf{x}_i, \mathbf{x}_{-i}) \geq \Pi_i(\tilde{\mathbf{x}}_i, \mathbf{x}_{-i}) \qquad \forall \tilde{\mathbf{x}}_i \in S_i,\ \forall i \in I \tag{5.2}$$

However, such pure Nash equilibria do not provably exist for general IPGs (Carvalho et al., 2022). For IPGs without pure Nash equilibria, common solution concepts include identifying approximate Nash equilibria or mixed equilibria. In an approximate Nash equilibrium, no player can benefit by more than ϵ from unilateral deviation (see Daskalakis et al., 2006; Papadimitriou, 2007):

$$\Pi_i(\mathbf{x}_i, \mathbf{x}_{-i}) + \epsilon \geq \Pi_i(\tilde{\mathbf{x}}_i, \mathbf{x}_{-i}) \qquad \forall \tilde{\mathbf{x}}_i \in S_i,\ \forall i \in I \tag{5.3}$$

A mixed equilibrium implies that players randomize over multiple strategies in the support of the equilibrium. The interpretation of such a randomized behavior can be difficult for some applications, particularly if modeled as single-shot games (Friedman and Zhao, 2021). In contrast, the interpretation of ϵ in an approximate Nash equilibrium is straightforward, be it the representation of partial player irrationality, players acting on imperfect input data, or simply unwillingness to change the status quo unless potential savings exceed a threshold. We will therefore rely on approximate Nash equilibria in the following.

Among other applications, IPGs, as introduced above, are used to model competitive location decisions in simultaneous move games (see, e.g., Chapter 4 or Crönert and Minner, 2021b). Methods to solve IPGs and closely related problem settings through the identification or selection of a Nash equilibrium are discussed by Sagratella (2019), Crönert and Minner (2021a, cf. Chapter 3) and Carvalho et al. (2022). In our motivating example of competitive retail location, such a Nash equilibrium represents location decisions of incumbent retail chains. In contrast, we are interested in the inverse problem, i.e., given a set of observed location decisions in an equilibrium, can we deduce hidden information on parameters that explain the observed choices?

5.2.2 Our context

We investigate a problem setting that involves three groups of actors: Customers (Section 5.2.3), incumbent retailers (Section 5.2.4) and a new entrant (Section 5.2.5). Appendix 5.D provides an overview of the notation used for the subproblems of the respective actors. Figure 5.1 shows the relationship and inter-dependencies between these actors. Incumbent retailers choose to open a subset of their candidate locations (red downward arrows) based upon their prior knowledge of each other's payoff structures and customer choice behavior (full information). They obtained this information by leveraging market experience, e.g., from historical sales data. The customers react to the chosen locations of all incumbent retailers by selecting certain stores they frequent (green upward arrows). We assume that the incumbents' selection of stores is in an (approximate) Nash equilibrium, i.e., is (near-)optimal given the customers' choice. A new entrant (the eye) observes the locations of the incumbents (blue horizontal arrows) and uses the gathered information to estimate the customers' choice parameters (blue dashed horizontal arrows). Subsequently, the new entrant can leverage the information to optimize their own location decision (not shown).

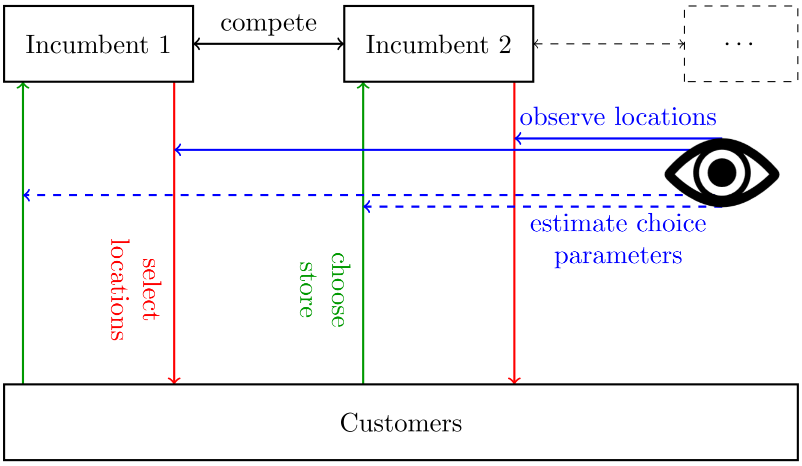

Figure 5.1: Problem structure and subproblems

While Figure 5.1 shows only a single observation, we assume the new entrant has access to multiple observations $o \in O$. From a theoretical perspective, the availability of multiple observations could be motivated through multiple equilibria in the same problem setting, or, in a more applied sense, the new entrant could observe multiple (sub-)regions. We implicitly assume that the observations (or the underlying regions) are sufficiently similar, i.e., customers behave similarly across all observations. The observations are limited to locations of incumbents in the same market segment, ignoring locations of retailers in other segments. While all retailers fiercely compete within the same segment, inter-segment effects are often negligible (see, e.g., Cleeren et al., 2010, for inter-segment competition of grocery retailers).

5.2.3 Customer choice

Customer patronage is driven by customers maximizing the utility they expect from the respective stores. The utility u^o_{ijk} that customers in location $j \in J$ receive from a facility $k \in K$ operated by an (incumbent) retail chain $i \in I$ in observation o is defined by a linear additive utility function. Note that this assumption is not as restrictive as it seems. With additional computational effort, Multiplicative competitive interaction (MCI) models (including the Huff-gravity model, Huff, 1964) or Multinomial logit choice (MNL) models could be used after reformulation into log-linear functions which are linear in the customer choice parameters. Yet, many studies on the predictive accuracy of market share models report little benefit of the more complex MCI or MNL models over linear models (Cooper and Nakanishi, 1988). The utility function builds on the following observations: The accessibility (e.g., proximity) of a store is a key driver for store choice, especially in grocery retailing (Statista, 2019a,b). Other important decision factors, such as differences in pricing, assortment, or product quality, depend on the store brand rather than the store location and are captured by an aggregate measure of retail chain brand attractiveness β_i. The vector $\boldsymbol{\beta} = (\beta_i)_{i \in I}$ summarizes β_i across all incumbents. The convenience of a location captures synergies with other points of interest in close proximity to the store (e.g., subway stations, gas stations, pharmacies, or organic supermarkets) in the spirit of multi-purpose shopping trips (see, e.g., Marianov et al., 2018). The customer choice parameter α describes the relative importance of accessibility over the convenience of multi-purpose shopping for customers. To ensure comparability, both accessibility and convenience are normalized between 0 and 1. We normalize the convenience measure g^o_k of store k in observation o, by comparing it with the the maximum convenience score across all store locations $\bar{g} = \max_{k \in K, o \in O}(g^o_k)$, yielding $\tilde{g}^o_k = \frac{g^o_k}{\bar{g}}$. Similarly, the normalized accessibility $\tilde{d}^o_{jk}$ is determined by the distance d^o_{jk} between store k and customers in j in relation to a maximum willingness to travel $\bar{d}$: $\tilde{d}^o_{jk} = \frac{\bar{d} - d^o_{jk}}{\bar{d}}$. I.e., a store k is "more accessible" to customer j the closer $\tilde{d}^o_{jk}$ is to 1. This definition of accessibility relates to research on the limited willingness of customers to travel beyond a threshold distance (Access Development, 2016; Anders, 2015). Other accessibility criteria, such as the driving time by car, taking into account road networks and congestion, or the availability of sufficient parking space, could also be included in alternative definitions of accessibility.

We assume that both store brand attractivity $0 \leq \beta_i \leq 1$ and the relative importance $0 \leq \alpha \leq 1$ of accessibility ($\tilde{d}^o_{jk}$) over convenience ($\tilde{g}^o_k$) are chosen from their respective discrete sets $\beta_i \in B$ and $\alpha \in A$. Jointly, they define the utility u^o_{ijk} as

$$u^o_{ijk} = \beta_i + \alpha \tilde{d}^o_{jk} + (1 - \alpha)\tilde{g}^o_k. \tag{5.4}$$

We assume that α and β_i ($i \in I$) are homogeneous across the examined market and are not location-dependent. This assumption is reasonable, as we examine competition within one segment (e.g., discount supermarkets). While the size of the customer base for this segment (i.e., the number of customers with a general interest in buying from any of the retailers) might differ between customer locations j according to their average income, the preferences $(\alpha, \boldsymbol{\beta})$ of customers willing to shop in the respective segment are unlikely to vary greatly between locations.

Based on the utility u^o_{ijk} (5.4) a store k of retailer i provides to customers located in j, the expected fraction f_{ij} of customers that patronize i's stores is assumed to follow Luce's choice axiom (see Luce, 1959)

$$f^o_{ij} = \frac{\sum\limits_{k \in K | d_{jk} \leq \bar{d}} x^o_{ik} u^o_{ijk}}{\sum\limits_{\tilde{i} \in I} \sum\limits_{k \in K | d_{jk} \leq \bar{d}} x^o_{\tilde{i}k} u^o_{\tilde{i}jk}} \qquad \forall i \in I,\ \forall j \in J,\ \forall o \in O. \tag{5.5}$$

This fractional patronage is based on the cumulative utility provided by all stores of i to customers at j, relative to the cumulative utility provided by all stores (including i's competitors) considered by customers at j.

Figure 5.2 illustrates a simple network of customer locations and potential store locations. Highlighting the interplay between arc-based parameters (d_{jk}), store-based parameters (g_k) and retail chain brand-based valuations (β_i). Here, the red chain (r) has two potential locations A and B, whereas the blue chain (b) has three potential locations C, D, and E. The dashed connection to E signifies that E is not reachable within the customers' maximum willingness to travel, i.e., $d_{j\mathrm{E}} > \bar{d}$. This means customers at j are willing to travel to 4 out of 5 potential locations.

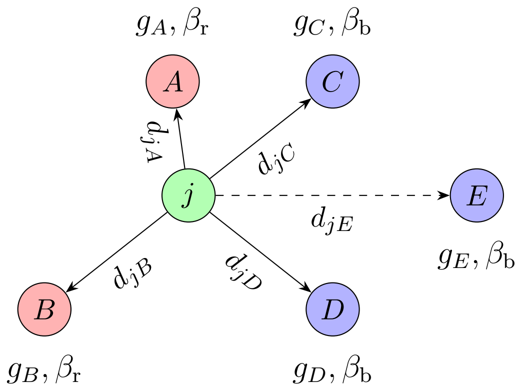

Figure 5.2: Network structure with one customer location ($|J|$=1), five potential retail locations ($|K| = 5$) and two incumbents ($I = \{\text{red (r)}, \text{blue (b)}\}$)

5.2.4 Forward problem: Incumbents' simultaneous location selection problem

We now describe the Integer programming game (IPG) between competing incumbents selecting their store locations simultaneously. Each incumbent retail chain i maximizes their own profit. For convenience, we assume costs and margins are annualized. For each observation o, the parameter $\hat{x}^o_{ik} = 1$ indicates retail chain i operates at location k. The variable $x^o_{ik} \in \{0,1\}$ represents an alternative location decision available to retail chain i. The profit $\Pi^o_i\left(\mathbf{x}^o_i, \hat{\mathbf{x}}^o_{-i}\right)$ retail chain i can attain with strategy $\mathbf{x}^o_i := (x^o_{ik})_{k\in K} \in S_i$ when other retail chains choose the observed strategies $\hat{\mathbf{x}}^o_{-i}$ during observation o, is given by

$$\Pi^o_i\left(\mathbf{x}^o_i, \hat{\mathbf{x}}^o_{-i}\right) = \underbrace{\sum_{j\in J} m^o_{ij} p^o_j f^o_{ij}}_{\text{annualized operating margin}} - \underbrace{\sum_{k\in K} c^o_{ik} x^o_{ik}}_{\text{annualized fixed costs}}. \tag{5.6}$$

The parameter m^o_{ij} is the average annual cumulative contribution margin for retail chain i per customer who resides in j, p^o_j is the (exogenous) population count or size of the customer base in j, and the (exogenous) parameter c^o_{ik} reflects all location-dependent costs, such as rental costs, and location-independent costs, such as salaries or equipment, for location k and retail chain i. Customer counts, contribution margins, and costs may differ between observations and are hence indexed with o.

Retail chains optimize their respective (annualized) profits, taking into consideration their competitors' location choices and the customer base and cost structure of the respective observation. Thus, for any given set of parameters $(\alpha, \boldsymbol{\beta})$ the strategies $\mathbf{x}^o_i$ of all retail chains $i \in I$ form an (approximate) Nash equilibrium (see Section 5.2.1) where the equilibrium strategies $\mathbf{x}_i^o$ depend on $(\alpha, \boldsymbol{\beta})$.

5.2.5 Inverse problem: The new entrant's parameter estimation problem

Upon observing the equilibrium store locations selected by incumbents $\hat{\mathbf{x}}_i^o$ (but not α, $\boldsymbol{\beta}$), the new entrant aims to estimate $(\alpha, \boldsymbol{\beta})$ such that they can use the estimated parameters to optimally select its own locations. This inverse problem can be described as follows. The new entrant is aware of the dependency $\hat{\mathbf{x}}_i^o(\alpha, \boldsymbol{\beta})$ between incumbent decision and customer choice parameters based on the incumbents' profit function. Estimating $(\alpha, \boldsymbol{\beta})$ from $\hat{\mathbf{x}}_i^o$ involves the inverse problem to identify a set of parameters $(\alpha, \boldsymbol{\beta})$ that best explains the currently observed structure of locations $\hat{\mathbf{x}}_i^o$ across the competing incumbent retail chains $i \in I$ as the outcome of the forward problem (5.4)-(5.6). Our inverse problem minimizes the L_1-norm $||\cdot||$ of $\boldsymbol{\delta} := (\delta_i^o)_{i \in I, o \in O}$, with δ_i^o as the unilateral improvement potential of retail chain i in observation o, i.e.,

$$\delta_i^o = \max_{\mathbf{x}_i^o \in S_i} \left\{ \Pi_i^o \left(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o \right) \right\} - \Pi_i^o \left(\hat{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o \right) \qquad \forall i \in I,\ \forall o \in O \tag{5.7}$$

$$\text{inverse IPG} : \min_{\alpha, \boldsymbol{\beta}} ||\boldsymbol{\delta}|| = \min_{\alpha, \boldsymbol{\beta}} \left|\left| \left(\max_{\mathbf{x}_i^o \in S_i} \left\{ \Pi_i^o \left(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o \right) \right\} - \Pi_i^o \left(\hat{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o \right) \right)_{i \in I, o \in O} \right|\right| \tag{5.8}$$

subject to (5.4)-(5.6). If an exact Nash equilibrium exists, the unilateral improvement potential across all retail chains and observations will be zero ($\min_{\alpha, \boldsymbol{\beta}} ||\boldsymbol{\delta}|| = 0$). In contrast, in an ϵ-Nash equilibrium, a (small) unilateral improvement potential remains for one ore more chains: $\epsilon = \min_{\alpha, \boldsymbol{\beta}} \max \boldsymbol{\delta}$. While the game theoretic literature commonly minimizes this maximum improvement potential ϵ across all retail chains (see Daskalakis et al., 2006; Papadimitriou, 2007); with the L_1-norm, we choose to minimize the cumulative deviations over all observations and retail chains, i.e., $\epsilon \approx \min_{\alpha, \boldsymbol{\beta}} ||\boldsymbol{\delta}||$. Thus, by equally weighing all observations, we avoid putting too much focus on a single observation outlier. Under the premise that the currently observed infrastructure is near-optimal, the parameter set solving (5.8) reflects customer decisions in practice: No retailer would have (substantial) incentive to deviate from the observed structure. We discuss an alternative to the L_1-norm based objective function (5.8) of this inverse problem in Appendix 5.E. Note that it is common in inverse optimization literature to continue to refer to parameters of the forward problem (here: $\alpha, \boldsymbol{\beta}$) as parameters, although they are represented in the inverse problem as decision variables. For convenience and to avoid confusion, we list all decision variables and (input) parameters in Appendix 5.D (Table 5.2).

5.3 Inverse optimization for parameter estimation in integer programming games

We motivate the need for a dedicated solution approach for inverse IPGs given in (5.8) based on a stylized example in Section 5.3.1, before Section 5.3.2 details our methodology. Sections 5.3.3 and 5.3.4 propose two improvements to the introduced methodology; tackling the existence of multiple solutions (parameter combinations) of the inverse IPG and the long runtimes in the forward problems, respectively.

5.3.1 A stylized example

Using a stylized example, we show that the enumeration of all constraints of the inverse problem is inefficient, the constraints can be non-convex, and that multiple solutions (parameter combinations) to the inverse IPG may exist, potentially in disconnected optimal regions.

Revisiting the example from Figure 5.2, two incumbent retail chains $i \in I = \{\text{red (r)}, \text{blue (b)}\}$, compete over customers located in j, i.e., $|J| = 1$. In contrast to Figure 5.2, we now ignore location E since customers are not willing to travel that far. To simplify exposition, we focus on a single observation, and thus drop all indices $o \in O$. Within this stylized example, retailers maximize a simplified payoff function where $p_j = 1$ and margins across customers and fixed costs across location are constant ($m = m_{ij}, c = c_{ik}$)

$$\Pi_i(\mathbf{x}_i, \mathbf{x}_{-i}) = m \sum_{j \in J} f_{ij} - c \sum_{k \in K} x_{ik} \tag{5.9}$$

with f_{ij} as given by (5.5).

We denote with $\mathbf{x}_\mathrm{r} = (x_\mathrm{rA}, x_\mathrm{rB})$ the decision of the red chain and with $\mathbf{x}_\mathrm{b} = (x_\mathrm{bC}, x_\mathrm{bD})$ the decision of the blue chain. Assume that the new entrant observes an (approximate) Nash equilibrium $(\hat{\mathbf{x}}_\mathrm{r}, \hat{\mathbf{x}}_\mathrm{b})$ where $\hat{\mathbf{x}}_\mathrm{r} = (x_\mathrm{rA} = 0, x_\mathrm{rB} = 1)$ (i.e., the red chain opens only location B) and $\hat{\mathbf{x}}_\mathrm{b} = (x_\mathrm{bC} = 1, x_\mathrm{bD} = 1)$ (i.e., the blue chain opens both locations C and D), as shown in Figure 5.3b. As we are dealing with discrete decisions, we cannot

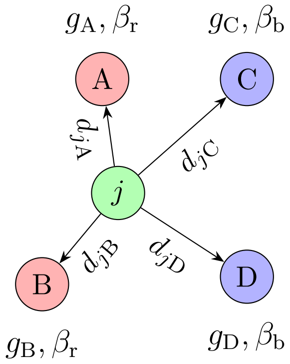

(a) Network with a single customer location j and a total of 4 potential locations

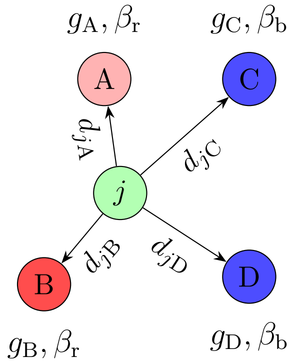

(b) Selected locations (dark red or dark blue, respectively) in an observed equilibrium

Figure 5.3: Network model with a single customer and two retail chains simultaneously selecting among four locations

exploit KKT conditions to characterize this equilibrium. Instead, we determine customer choice parameters by ensuring optimality of the observed decisions by comparing with all possible unilateral deviations

$$\max_{\mathbf{x}_i \in S_i} \Pi_i\left(\mathbf{x}_i, \hat{\mathbf{x}}_{-i}\right) \geq \Pi_i\left(\mathbf{x}_i, \hat{\mathbf{x}}_{-i}\right) \qquad \forall i \in \{\mathrm{r}, \mathrm{b}\},\ \forall \mathbf{x}_i \in S_i, \tag{5.10}$$

yielding

$$\delta_i \geq \Pi_i\left(\mathbf{x}_i, \hat{\mathbf{x}}_{-i}\right) - \Pi_i\left(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}_{-i}\right) \qquad \forall i \in \{\mathrm{r}, \mathrm{b}\},\ \forall \mathbf{x}_i \in S_i. \tag{5.11}$$

We assume an exemplary set of parameters ($\tilde{d}_{\mathrm{A}} = 0.5,\ \tilde{d}_{\mathrm{B}} = 0.4,\ \tilde{d}_{\mathrm{C}} = 0.7,\ \tilde{d}_{\mathrm{D}} = 0.4,\ \tilde{g}_{\mathrm{A}} = 0.4,\ \tilde{g}_{\mathrm{B}} = 0.5,\ \tilde{g}_{\mathrm{C}} = 0.45,\ \tilde{g}_{\mathrm{D}} = 0.65,\ m = 1000,\ c = 160$) to explore some key problem characteristics in the following. Because $\hat{\mathbf{x}}_{\mathrm{r}} = (x_{\mathrm{r}A} = 0, x_{\mathrm{r}B} = 1)$ and $\hat{\mathbf{x}}_{\mathrm{b}} = (x_{\mathrm{r}C} = 1, x_{\mathrm{r}D} = 1)$, we can use (5.6) to show that red's payoff in the observed solution is

$$\begin{aligned} \Pi_{\mathrm{r}}\left(\hat{\mathbf{x}}_{\mathrm{r}}, \hat{\mathbf{x}}_{-\mathrm{r}}\right) &= 1000 \sum_{j \in J} f_{\mathrm{r}j} - 160(\hat{x}_{\mathrm{r}A} + \hat{x}_{\mathrm{r}B}) \\ &= \left(1000 \frac{\beta_{\mathrm{r}} + 0.4\alpha + 0.5\left(1 - \alpha\right)}{\beta_{\mathrm{r}} + 2\beta_{\mathrm{b}} + 1.5\alpha + 1.6\left(1 - \alpha\right)} - 160\right). \end{aligned} \tag{5.12}$$

Also, we can consider an alternative solution for red that constitutes a unilateral move so that $\mathbf{x}_\mathrm{r} = (1,1)$ and $\hat{\mathbf{x}}_\mathrm{b} = (1,1)$ and the corresponding payoff is

$$\Pi_r\left(\mathbf{x}_\mathrm{r}, \hat{\mathbf{x}}_{-\mathrm{r}}\right) = 1000 \sum_{j \in J} f_{\mathrm{r}j} - 160(x_{\mathrm{r}A} + x_{\mathrm{r}B})$$
$$= \left(1000 \frac{2\beta_\mathrm{r} + 0.9\alpha + 0.9\,(1-\alpha)}{2\beta_\mathrm{r} + 2\beta_\mathrm{b} + 2\alpha + 2\,(1-\alpha)} - 360\right). \quad (5.13)$$

When comparing red's payoff in the observed solution $\Pi_\mathrm{r}\left(\hat{\mathbf{x}}_\mathrm{r}, \hat{\mathbf{x}}_{-\mathrm{r}}\right)$, with the payoff $\Pi_\mathrm{r}\left(\mathbf{x}_\mathrm{r}, \hat{\mathbf{x}}_{-\mathrm{r}}\right)$ in an alternative solution where red opens both locations $\mathbf{x}_\mathrm{r} = (1,1)$, one of the constraints in (5.11) yields

$$\begin{aligned} \delta_r \geq & \left(1000 \frac{2\beta_\mathrm{r} + 0.9\alpha + 0.9\,(1-\alpha)}{2\beta_\mathrm{r} + 2\beta_\mathrm{b} + 2\alpha + 2\,(1-\alpha)} - 360\right) \\ & - \left(1000 \frac{\beta_\mathrm{r} + 0.4\alpha + 0.5\,(1-\alpha)}{\beta_\mathrm{r} + 2\beta_\mathrm{b} + 1.5\alpha + 1.6\,(1-\alpha)} - 160\right). \end{aligned} \quad (5.14)$$

The example of (5.14) shows that the constraints defining δ_i in (5.11) can be non-convex. Solving (5.8) (i.e., $\min_{\alpha,\boldsymbol{\beta}} ||\boldsymbol{\delta}||$) under constraints (5.11) defines the optimal region for the unknown parameters $(\alpha, \boldsymbol{\beta})$. To enable two-dimensional visualization, we enforce an additional constraint on β_i, imposing a cumulative brand effect: $\beta_\mathrm{r} + \beta_\mathrm{b} = 1$. This fixed relation of a customers appreciation of one chain over the other allows us to depict two dimensional optimal regions for two examples in Figure 5.4. We re-use the parameter set from above, but slightly increase location fixed costs c for Figure 5.4a compared with Figure 5.4b. In both examples, the equilibrium is exact (i.e., $\epsilon = \min_{\alpha,\boldsymbol{\beta}} ||\boldsymbol{\delta}|| = 0$). The resulting optimal regions have non-linear boundaries (Figure 5.4a). In some cases, they consist of multiple, disconnected and non-convex sub-regions (Figure 5.4b). As we move from higher to lower fixed costs, many combinations of $\alpha, \boldsymbol{\beta}$ that were optimal before (Figure 5.4a) are now excluded from the optimal region (Figure 5.4b). For these excluded combinations, the constraint (5.14) becomes binding: Due to the reduced fixed costs, a decision by red to open both locations $\mathbf{x}_\mathrm{r} = (1,1)$ is favored over the observed solution $\hat{\mathbf{x}}_\mathrm{r} = (0,1)$ in which only a location at B is opened ($x_{\mathrm{rB}} = 1$).

In this stylized example, enumerating all possible alternative strategies $\mathbf{x}_i \in S_i$ per player $i \in I$ leads to $\sum_{i \in I} |S_i| = 8$ constraints (5.11) defining $\boldsymbol{\delta}(\alpha, \boldsymbol{\beta})$. However, it is clear that in realistically-sized problem instances, such a full enumeration is intractable, and many of the enumerated constraints would be dominated. For example, in the case of Figure 5.4a, two out of the eight constraints suffice to fully describe the solution space

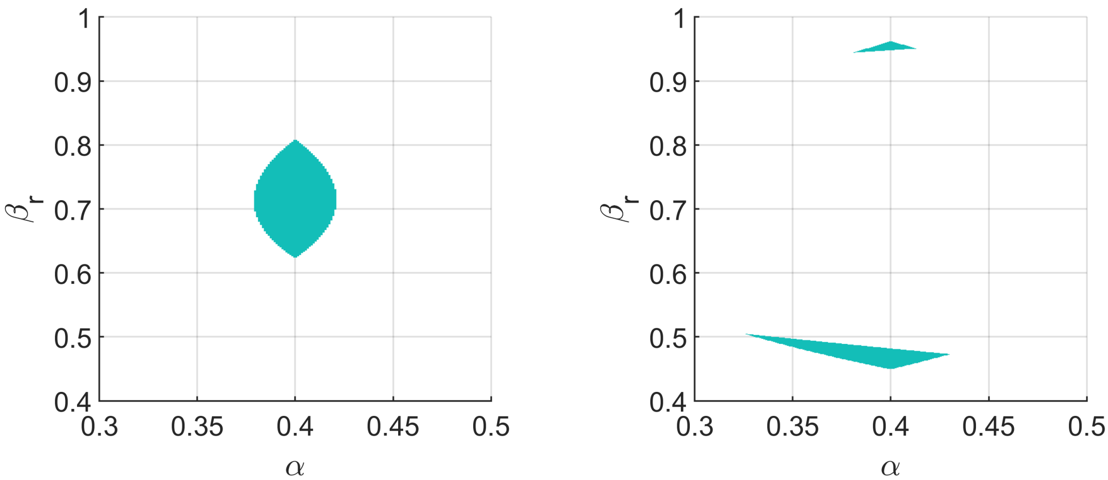

(a) Optimal region, assuming a relative chain attractiveness of $\beta_\mathrm{r} + \beta_\mathrm{b} = 1$ and fixed costs of $c = 170$

(b) Optimal region, assuming a relative chain attractiveness of $\beta_\mathrm{r} + \beta_\mathrm{b} = 1$ and fixed costs of $c = 160$

Figure 5.4: Optimal region ($||\delta|| = 0$) for the inverse network problem

for $||\boldsymbol{\delta}|| = 0$. Therefore, we propose a row-generation approach, called invIPG, limiting the number of enumerated constraints (5.11) while determining a feasible parameter combination that solves the respective inverse IPG.

5.3.2 Solution approach: invIPG

We build on the general concept from Wang (2009) and split the problem into a master problem, identifying the parameter combination $(\alpha, \boldsymbol{\beta})$, and row-generating subproblems, solving the forward problem for each player and each observation. However, in contrast to Wang (2009), we are not interested in finding solutions to the master problem that minimize deviations from a prior (e.g., an initial guess). Instead, we identify solutions that minimize the unilateral improvement potential $\boldsymbol{\delta}$ across all players. This allows us to identify parameter combinations that define an (approximate) equilibrium in an IPG.

Figure 5.5 summarizes the concept of the proposed invIPG method. We initialize the algorithm with an empty set of subproblem solutions, $\tilde{S}_i = \emptyset$, and solve the master problem which represents a relaxed version of the inverse problem: Given observations $\hat{\mathbf{x}}$, the inverse problem minimizes the unilateral improvement potential $\boldsymbol{\delta} := (\delta_i^o)_{i \in I, o \in O}$ between the observed solution $\hat{\mathbf{x}}_i^o$ and the optimal (for a given $\alpha, \boldsymbol{\beta}$) solutions $\mathbf{x}_i^o$ across all observations $o \in O$ and players $i \in I$, by choosing the estimation parameters $\alpha, \boldsymbol{\beta}$.

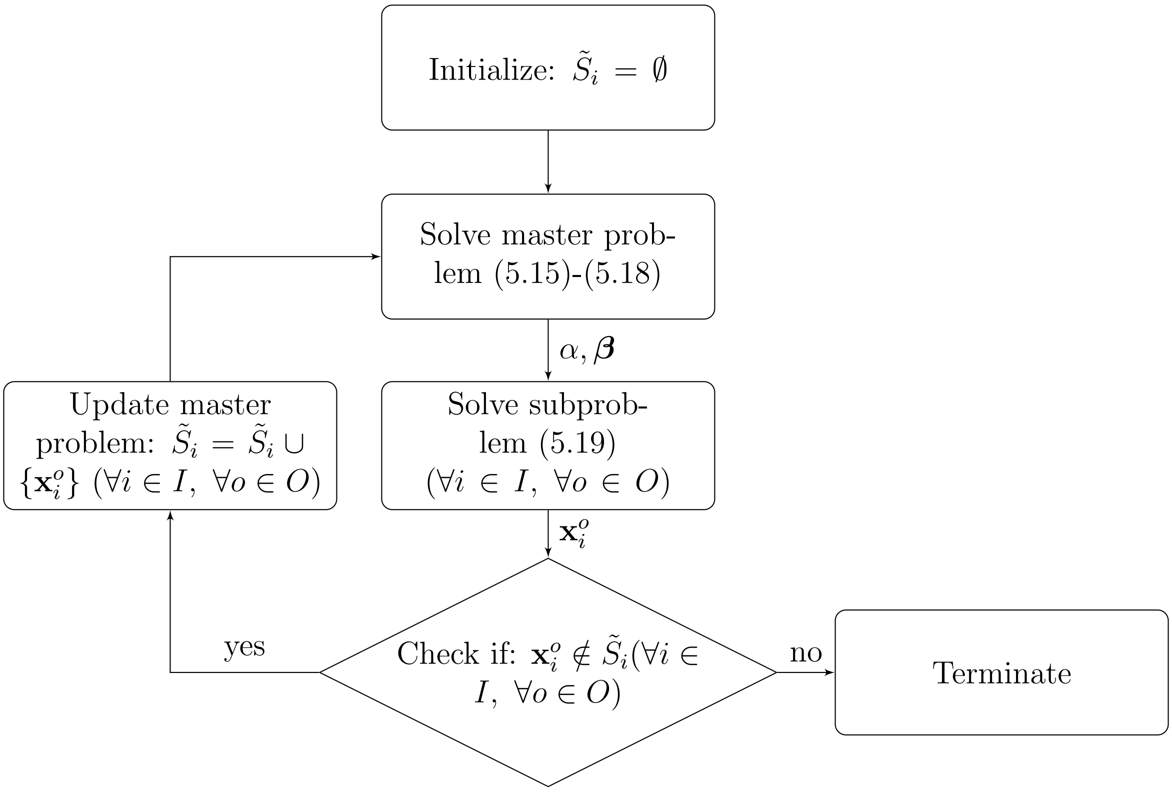

Figure 5.5: Flow chart of the invIPG algorithm

Through enumeration of all alternative strategies $\tilde{\mathbf{x}}_i \in S_i$ one can reformulate this inverse problem (5.8) to

Master problem: $$\min_{\alpha,\boldsymbol{\beta}} ||\boldsymbol{\delta}|| = \min_{\alpha,\boldsymbol{\beta}} ||(\delta_i^o)_{i\in I, o \in O}|| \tag{5.15}$$

$$\delta_i^o = \Pi_i^o\left(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o\right) - \Pi_i^o\left(\hat{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o\right) \qquad \forall i \in I,\ \forall o \in O \tag{5.16}$$

$$\Pi_i^o\left(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o\right) \geq \Pi_i^o\left(\tilde{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o\right) \qquad \forall i \in I,\ \forall o \in O,\ \tilde{\mathbf{x}}_i^o \in S_i. \tag{5.17}$$

Instead of ensuring that $\mathbf{x}_i^o$ are optimal compared to all alternative strategies $\tilde{\mathbf{x}}_i \in S_i$, we relax (5.16)-(5.17) to apply to the gradually increasing subset $\tilde{S}_i$ of S_i only ($\tilde{S}_i \subseteq S_i$)

$$\delta_i^o \geq \Pi_i^o\left(\tilde{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o\right) - \Pi_i^o\left(\hat{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o\right) \qquad \forall i \in I,\ \forall o \in O,\ \tilde{\mathbf{x}}_i^o \in \tilde{S}_i. \tag{5.18}$$

For the competitive retail location problem, the profit functions Π_i^o are defined in (5.4)-(5.6) of Section 5.2.4. In the described master problem, location decisions $(\tilde{\mathbf{x}}_i^o, \hat{\mathbf{x}}_i^o)$ are enumerated, but estimation parameters $(\alpha, \boldsymbol{\beta})$ represent decision variables. As the profit functions (5.4)-(5.6) are non-linear in the estimation parameters $(\alpha, \boldsymbol{\beta})$ we provide a linearization in Appendix 5.B.

After identifying an initial feasible parameter set through the master problem, we iteratively reconstruct the initially relaxed conditions (5.16)-(5.17) through (5.18) by increasing $\tilde{S}_i$ with solutions from the forward problem: Based on the current (incumbent) solution of the parameter set from the master problem, the forward problem (Section 5.2.4) is solved to optimality for every player and observation, returning an (integer) solution $\mathbf{x}_i^o$.

$$\text{Subproblem:} \quad \max_{\mathbf{x}_i^o} \Pi_i^o\left(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o\right) = \max_{\mathbf{x}_i^o} \sum_{j \in J} m_{ij}^o p_j^o f_{ij}^o - \sum_{k \in K} c_{ik}^o x_{ik}^o \quad \forall i \in I,\ \forall o \in O \tag{5.19}$$

Note that in this subproblem, $(\alpha, \boldsymbol{\beta})$ are parameters (current incumbent solution of the master problem) rather than decision variables. However, the profit functions (5.4)-(5.6) are now non-linear in the location decisions $\mathbf{x}_i^o$. This new non-linearity of the subproblem can be linearized as a product between a binary decision x_{ik}^o and a continuous decision f_{ij}^o (for reference, see Appendix 5.A). The solution $\mathbf{x}_i^o$ to this subproblem is added to $\tilde{S}_i$

$$\tilde{S}_i = \tilde{S}_i \cup \{\mathbf{x}_i^o\} \qquad \forall i \in I,\ \forall o \in O. \tag{5.20}$$

Any feasible parameter set in the master problem must now ensure (approximate) optimality of the observations when compared with the newly added solutions from the subproblems in (5.18). The increased master problem and subproblems are solved until the algorithm converges when the subproblem does not result in new cuts for the master problem (Theorem 5.1).

Theorem 5.1. *The invIPG converges in a finite number of iterations.*

Proof. S_i denotes the finite set of all possible strategies. For the competitive retail location problem, $\mathbf{x}_i^o \in S_i$ with $\mathbf{x}_i^o = (x_{ik}^o)_{k \in K}$. As the number of observations $|O|$, the number of potential locations $|K|$ and the number of retail chains $|I|$ is finite, this set S_i is finite. In each iteration, at least one strategy $\mathbf{x}_i^o$ is added to $\tilde{S}_i \subseteq S_i$. By consequence, in the worst case, the algorithm terminates in a finite number of iterations after full enumeration of all pure strategies $(\tilde{S}_i = S_i)$. □

Remark 5.1. *In Section 5.2.1, we allow action sets S_i to be defined by any (potentially non-linear) real-valued functions w_q but restrict ourselves to IPGs, in which all decision variables are integer and bounded, to ensure S_i is finite. Alternatively, one could allow mixed-integer IPGs in combination with linear inequalities to define S_i and rely on the finite number of extreme points of a polyhedron to prove convergence in a finite number of iterations in Theorem 5.1 (see Wang, 2009, for non-competitive inverse mixed-integer problems).*

Note that Theorem 5.1 does not make any assumptions with regard to the nature of the observed equilibria (ϵ-Nash or Nash equilibrium). In case all observations represent exact Nash equilibria, the algorithm terminates with $\epsilon = \min_{\alpha,\boldsymbol{\beta}} ||\boldsymbol{\delta}|| = 0$. In case any observation is an $\epsilon-$Nash equilibrium, the algorithm still converges and terminates with $\min_{\alpha,\boldsymbol{\beta}} ||\boldsymbol{\delta}|| > 0$.

5.3.3 Existence of multiple parameter solutions

An IPG can have multiple equilibria; the identification of the expected equilibria among the set of all equilibria is of considerable importance and an interesting field of research (equilibrium selection theory, see, e.g., Chapter 3 or Crönert and Minner, 2021a; Harsanyi, 1995). For the invIPG, the existence of multiple equilibria of the underlying IPG is not an issue – either the final (approximate) equilibrium can be readily observed in practice, or multiple equilibria can be used as multiple observations in the algorithm. Yet, as the stylized example in Section 5.3.1 shows, the solution to the invIPG itself may not be unique, with multiple parameter combinations satisfying (5.15). Egri et al. (2014) make a similar observation for their formulation of an inverse economic lot sizing problem and propose a bounding box of the feasible region of all solutions to their inverse optimization problem. In contrast to their problem formulation, the optimal region in our inverse problem may be non-convex and feasible regions might be unconnected (see Figure 5.4b in Section 5.3.1). Consequently, a single bounding box cannot adequately describe the feasible region for these instances.

The invIPG terminates with the identification of a single parameter combination $(\alpha, \boldsymbol{\beta})$ that minimizes the unilateral improvement potential $\boldsymbol{\delta}$ across all players and observations. Similar to the identification of all solutions to integer programs (see, e.g., Tsai et al., 2008), one could iteratively enumerate all possible solutions by excluding identified solutions through cuts. As both $\alpha \in A$ and $\boldsymbol{\beta} \in B$ are defined on a finite solution space (discrete and bounded), using the same line of argumentation as in Theorem 5.1 such an approach would terminate in a finite number of iterations.

5.3.4 Approximate cuts

The forward problem turns out to be hard to solve for larger problem instances. For most instances, the solver (here: Gurobi 9.5) quickly identifies the optimal solution and spends most computation time proving this optimality. Bodur et al. (2022) observe a similar behavior for non-competitive inverse mixed-integer optimization and propose to stop the cut-generating subproblem early in case a feasible solution has been found and the time threshold has been exceeded. Related approaches are also common in Benders decomposition approaches in the form of inexact Benders cuts (see, e.g., Zakeri et al., 2000). Instead of a threshold on the total solution time of the subproblem as suggested by Bodur et al. (2022), we prematurely stop the solver for the subproblem in case no improved solution has been found within a predefined timelimit (see extended algorithmic flow chart in the Appendix 5.C, Figure 5.9). This adapted stopping criterion avoids early termination when the solver continues to find better solutions, as observed in some preliminary numerical experiments. In combination with (5.18) and (5.20), approximate solutions represent valid and effective inequalities which are added to the master problem.

Lemma 5.1. *Approximately optimal solutions to the subproblem are valid and effective inequalities for the relaxed master problem.*

Proof. Let $\tilde{\mathbf{x}}_i^o$ be an approximate solution to the subproblem (5.19). I.e., there exists a solution $\mathbf{x}_i^o \in S_i^o \setminus \tilde{S}_i^o$, such that $\Pi_i^o(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o) \geq \Pi_i^o(\tilde{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o)$ for the current $(\alpha, \boldsymbol{\beta})$. However the solution $\tilde{\mathbf{x}}_i^o$ ensures higher payoffs for i, compared with all already enumerated solutions $\check{\mathbf{x}}_i^o \in \tilde{S}_i^o$: $\Pi_i^o(\tilde{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o) > \Pi_i^o(\check{\mathbf{x}}_i^o, \hat{\mathbf{x}}_{-i}^o)$.

This solution represents a valid cut, since in the original inverse problem (5.18) has to hold for any $\tilde{\mathbf{x}}_i^o \in S_i$ (see (5.7)). It is effective, since it improves the players profit Π_i^o for the incumbent parameter combination $(\alpha, \boldsymbol{\beta})$ compared to all other already enumerated solutions $\bar{S}_i$. □

As before, we repeatedly solve master and subproblems until no new approximate cuts can be identified within the timelimit. However, in an additional final step, we solve the subproblem to optimality once for every observation $o \in O$ and retail chain $i \in I$ to identify any additional (exact) cuts. If this final step also does not yield any new cuts, the algorithm terminates; otherwise, it continues with identifying additional exact cuts. Even though cuts are approximated in early iterations, this final step ensures optimality of the identified parameter combination after termination. We assess the implications of this extension on runtime in a numerical study in Section 5.4.2.

5.4 Numerical study

We implement the invIPG in Python 3.8 and use Gurobi 9.5 as a solver. Reported runtimes are based on an 8 core CPU with a 2.1 GHz base frequency (3.7 GHz boost) and 22 GB RAM.

We examine problem instances with two incumbent retailers ($|I| = 2$) and a single new entrant. Parameter data $(d_{jk}, g_k, p^o_j, c^o_{ik})$ is common knowledge and sampled randomly. For population size p^o_j and costs c^o_{ik} we use a uniform distribution ($p^o_j, c^o_{ik} \sim U(45, 55)$).

Upper and lower bounds of these distributions are chosen to define a meaningful problem in which it is profitable for incumbents to open some but not all potential locations. A $\pm 10\%$ deviation of minimum and maximum values from the distribution mean allows for sufficiently diverse locations while avoiding trivial problems in which some locations are clearly optimal, regardless of the customer choice parameters $(\alpha, \boldsymbol{\beta})$.

To determine d_{jk}, we calculate the Euclidean distance between randomly sampled locations (J, K), uniformly distributed on a 100-by-100 plane. To generate g_k, we sample additional random locations of 50 complementary stores from the same distribution and determine the average proximity of a location k to these complementary stores. Profit margins are assumed to be constant; $m_{ij} = 1, \ \forall i \in I, \ \forall j \in J$. We examine two smaller problem instances ($|J| = |K| = 10, \ |J| = |K| = 20$) with a maximum customer willingness to travel $\bar{d} = 100$ and a larger instance ($|J| = |K| = 30$) with a reduced maximum distance between customers and the stores relevant to them ($\bar{d} = 50$).

The ground-truth values of the customer choice parameters $\hat{\alpha}$ and $\hat{\boldsymbol{\beta}}$ are known to the incumbents but unknown to the new entrant. In the first two experiments (Sections 5.4.1 and 5.4.2), we assume constant values $\hat{\alpha} = 0.4$ and the relative appreciation of customers for one retail chain over the other to be fixed ($\hat{\beta}_1 = 1.5 \cdot \hat{\beta}_0$). This limits random effects and reduces the number of required samples, enabling a ceteris paribus examination of the effects of an increase in the number of observations (Section 5.4.1) or the usage of approximate cuts (Section 5.4.2). For the last experiment (Section 5.4.3), which compares the decision-making of a new entrant applying the invIPG to benchmarks based on full information and scenario sampling approaches, we increase the number of samples and choose customer choice parameters from a uniform distribution ($\hat{\alpha}, \hat{\beta}_0, \hat{\beta}_1 \sim U(0,1)$) Equilibrium locations $\hat{\mathbf{x}}$ observed by the new entrant are generated from these ground-truth values through a series of iterated best responses between the two incumbents. Note that this approach does not provably lead to a pure-strategy equilibrium; instead, it may very well be the case that the iterated best responses cycle. In one-third of the examined instances, we observe this cycling behavior for one of the observations. In these cases, we break the cycle and let the current solution represent the observation of an approximate equilibrium. In addition to these approximate equilibria based on cycling best-responses, in experiments with noisy observations, we assume incumbents do not know $\hat{\alpha}, \hat{\beta}_i$ exactly during the generation of equilibrium observations, but instead rely on noisy distortions of $\hat{\alpha}$: $\tilde{\alpha} \sim N(\mu = \hat{\alpha}, \sigma = 0.05)$ and $\hat{\beta}_i$: $\tilde{\beta}_i \sim N(\mu = \hat{\beta}_i, \sigma = 0.05)$. These absolute distortions of $\sigma = 0.05$ correspond to a relative coefficient of variation of $cv = 10\%$ and could represent unobserved parameters or simply imperfections in the incumbents' sales data and their estimation of customer attraction parameters. The new entrant estimates α and $\boldsymbol{\beta}$ from their respective discrete sets A and B. In the absence of additional information, we assume sets A and B with equal spacing and a granularity of two decimal points, i.e., $A = B = \{0, 0.01, 0.02, \ldots, 1\}$.

5.4.1 Existence of multiple solutions (parameter combinations)

We discuss the possible existence of multiple parameter combinations that solve the inverse IPG in Section 5.3.1 and 5.3.3. Before applying the invIPG to the competitive retail location problem, we first assess the characteristics that drive the existence of multiple solutions and the likelihood of multiple solutions in different problem settings. This experiment illustrates the conclusiveness of observations, determining the extent to which the incumbent's location decisions contain descriptive information on customer choice parameters – or whether arbitrary parameters would lead to the same observations. We investigate the influence of the number of observations ($|O| \in \{4, 6, 8\}$) and the noisiness of the observations ($\sigma \in \{0, 0.05\}$) on the (relative) size of the optimal region Γ (i.e., the fraction of all possible parameter combinations that equally well explain the observations). Define γ as a binary indicator determining whether a parameter set $(\alpha, \boldsymbol{\beta})$ is optimal, i.e., $\gamma(\alpha, \boldsymbol{\beta}) = 1$, iff $||\boldsymbol{\delta}(\alpha, \boldsymbol{\beta})|| = \min_{\tilde{\alpha}, \tilde{\boldsymbol{\beta}}} ||\boldsymbol{\delta}(\tilde{\alpha}, \tilde{\boldsymbol{\beta}})||$.

The relative size of the optimal region Γ is the area of parameter combinations that lead to a minimal value of $\boldsymbol{\delta}$ ($\gamma = 1$) compared to the size of the full solution space, i.e., $\Gamma = \frac{\int_\alpha \int_{\boldsymbol{\beta}} \gamma(\alpha, \boldsymbol{\beta}) d\alpha d\boldsymbol{\beta}}{\int_\alpha \int_{\boldsymbol{\beta}} 1 d\alpha d\boldsymbol{\beta}}$. Note that the objective of the invIPG was the identification of a single parameter combination satisfying $\gamma(\alpha, \boldsymbol{\beta}) = 1$, rather than all optimal parameter combinations. Therefore, it cannot be applied to calculate Γ. In the absence of an analytical solution to Γ, we enumerate all possible combinations of $(\alpha, \boldsymbol{\beta})$ and calculate the respective $\gamma(\alpha, \boldsymbol{\beta})$. For all of these combinations, this implies solving a mixed-integer program (5.7) for each $o \in O$ and $i \in I$. With a granularity of two decimal points ($|A| = |B| = 100$) for α, β_0, β_1 we need to enumerate and solve $100^3 = 1$ million different parameter combinations. It is evident that such an approach would not be feasible in an experimental design involving a large number of random samples. Therefore, we examine only parts of the solution space by introducing an additional constraint for the cumulative brand effect: $\beta_0 + \beta_1 = \mathcal{B}$. We examine three different values for $\mathcal{B} \in \{0.5, 0.75, 1\}$. This limitation of the naïve enumerative approach demonstrates the benefit of the invIPG, which requires no such assumption.

Table 5.1 examines the number of disconnected optimal regions, the cumulative size of these regions (Γ) and the minimum error (L_2-norm) between a solution within any of the optimal regions and the original ground-truth parameter values $\hat{\alpha}, \hat{\boldsymbol{\beta}}$. Each row summarizes results for $n = 20$ samples, aggregating samples across different values of $\mathcal{B}$ or $|O|$, respectively. In all but 4% of samples, the problems yield only one optimal region. Both a higher number of observations and the existence of noise effectively reduce the size of the optimal region(s) (Γ). However, noisy observations clearly hinder solution quality as the error between optimal region and ground-truth value increases significantly in presence of noise. Even in the absence of noise, some of the observations represent approximate equilibria. These approximate equilibria arise if the iterated best responses used to generate observations do cycle. For these approximate equilibria, the ground-truth value is not necessarily part of the optimal region, leading to outliers with an error greater than 0.

5.4.2 Effects of approximate cuts

We now move from the complete enumeration of the solution space to the application of the invIPG for larger problem instances. Section 5.3.4 discusses the possibility of further speeding up the runtime of the invIPG through the usage of approximate instead of exact cuts. This experiment examines the advantage of approximate cuts on solution time. In our implementation, the incumbent subproblem solution terminates with an approximate cut in case no improved solution has been found within a timelimit of 200 seconds. Each boxplot in Figure 5.6 summarizes runtimes for $n = 20$ samples of a larger problem ($|J| = |K| = 30$, $|O| = 6$, $\hat{\alpha} = 0.4$, $\hat{\beta}_0 = 0.4$, $\hat{\beta}_1 = 0.6$). Lower and upper box limits show the first and third quartile, the green horizontal line indicates the median, and whiskers

Table 5.1: Solution characteristics for $|J| = |K| = 10$ in the absence ($cv = 0\%$) and presence ($cv = 10\%$) of noise, $n = 20$ samples per row

			Number of optimal regions			Size of optimal regions			Error $\|\|(\alpha, \boldsymbol{\beta}) - (\hat{\alpha}, \hat{\boldsymbol{\beta}})\|\|_2$		
$cv(\%)$	O	$\mathcal{B}$	min	mean	max	min	mean	max	min	mean	max
0	4	0.5	1	1.0	1	0.0	0.03	0.17	0.0	0.0	0.06
0	4	0.75	1	1.0	1	0.0	0.02	0.17	0.0	0.02	0.13
0	4	1.0	1	1.15	3	0.0	0.01	0.14	0.0	0.04	0.18
0	6	0.5	1	1.0	1	0.0	0.01	0.05	0.0	0.01	0.06
0	6	0.75	1	1.0	1	0.0	0.01	0.06	0.0	0.02	0.13
0	6	1.0	1	1.1	2	0.0	0.0	0.04	0.0	0.03	0.15
0	8	0.5	1	1.0	1	0.0	0.01	0.05	0.0	0.01	0.06
0	8	0.75	1	1.05	2	0.0	0.0	0.02	0.0	0.03	0.23
0	8	1.0	1	1.1	2	0.0	0.0	0.02	0.0	0.04	0.15
10	4	0.5	1	1.0	1	0.0	0.02	0.1	0.0	0.02	0.06
10	4	0.75	1	1.05	2	0.0	0.01	0.07	0.0	0.03	0.1
10	4	1.0	1	1.0	1	0.0	0.01	0.05	0.0	0.04	0.18
10	6	0.5	1	1.0	1	0.0	0.01	0.05	0.0	0.02	0.06
10	6	0.75	1	1.05	2	0.0	0.01	0.06	0.0	0.02	0.05
10	6	1.0	1	1.05	2	0.0	0.01	0.04	0.0	0.04	0.2
10	8	0.5	1	1.11	3	0.0	0.0	0.05	0.0	0.04	0.13
10	8	0.75	1	1.05	2	0.0	0.0	0.02	0.0	0.04	0.15
10	8	1.0	1	1.05	2	0.0	0.0	0.02	0.0	0.06	0.15

extend to largest or smallest data points within a maximum distance of 1.5 times the interquartile range from the box limits. The illustration focuses on runtimes, as the identified solutions are identical for both approaches. Using approximate cuts reduces the median runtime by 38% from $\approx$ 5.7 hours to $\approx$ 4.3 hours. It proves particularly useful for hard-to-solve instances with very large runtimes in the standard approach. In all cases, both approaches require the same number of cuts and the same number of iterations. In fact, the identified approximate cuts turned out to be identical to the exact cuts in the standard approach, although this is not necessarily the case in general.

5.4.3 Comparison to full information and sampling-based benchmarks

To quantify the advantage of the invIPG for a new entrant, we assume the entrant solves a sequential competitive facility location problem for given incumbent locations. We determine profits of the entrant's location decision based on inversely estimated parameters (invIPG) and compare it to new entrants' decisions based on information

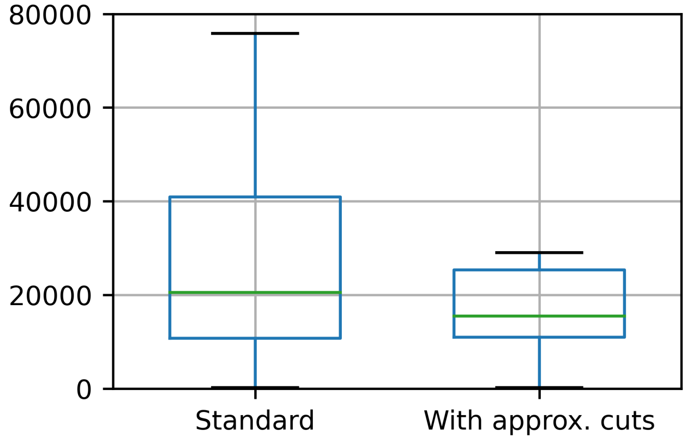

Figure 5.6: Runtime (in seconds) of the invIPG in its standard version and with approximate cuts for $n = 20$ samples per boxplot, not shown are three instances exceeding the timelimit of 80k seconds for both approaches

regarding the distribution from which the parameters are drawn. For the latter, we distinguish between an approach relying on the Distribution mean (DM) and an approach representing the distribution through Sample average approximation (SAA). All approaches (invIPG, DM, SAA) are compared to decision-making under full information about the ground-truth parameters $(\hat{\alpha}, \hat{\boldsymbol{\beta}})$ used to generate the observed equilibria.

A new entrant who bases their store location decision on information about the Distribution mean (DM) of these parameters would assume $\alpha = \mu_{\hat{\alpha}} = 0.5$, $\beta_0 = \mu_{\hat{\beta}_0} = 0.5$, and $\beta_1 = \mu_{\hat{\beta}_1} = 0.5$. For the SAA, we assume that the new entrant uses 64 scenarios, with all possible combinations of α, β_0, β_1 being chosen from $\{0.2, 0.4, 0.6, 0.8\}$. Since the ground-truth values of $\hat{\alpha}, \hat{\beta}_0, \hat{\beta}_1$ are assumed to be uniformly distributed, this scenario approach enables a comprehensive and reproducible representation of the underlying distribution and was thus chosen over other sampling techniques (e.g., Monte Carlo method, Latin hypercube sampling). All scenarios are assigned equal probabilities and the new entrant selects locations yielding the highest average profit across all scenarios (for a detailed introduction to SAA, see Birge and Louveaux, 2011). In contrast, a new entrant relying on the invIPG uses all available observations in O to estimate these parameters. In addition, for all approaches (invIPG, DM, SAA) the new entrant needs an estimate β_2 of the customers' appreciation of their own brand. As we are interested in the value of information extracted from the observed location structure of incumbents, we assume this estimate equals the ground-truth value $\hat{\beta}_2 = \hat{\beta}_2 \sim U(0, 0.5)$ in all cases. This distribution of $\hat{\beta}_2$ has a lower maximum value compared to $\hat{\beta}_0$ and $\hat{\beta}_1$ as we expect a reduced initial customer appreciation for the new entrant. We identify the entrant's optimal location decisions based on the three different approaches and determine the respective profits of these decisions under the ground-truth customer attraction parameters. Figure 5.7 and 5.8 show the relative profit, normalized to a profit of 100% under full information, for two problem sizes, with $n = 100$ (for $|J| = |K| = 10$) and $n = 40$ (for $|J| = |K| = 20$) samples, respectively. We employ a timelimit of 20k seconds for the invIPG and use the current incumbent solution in case the invIPG did not converge within the timelimit.

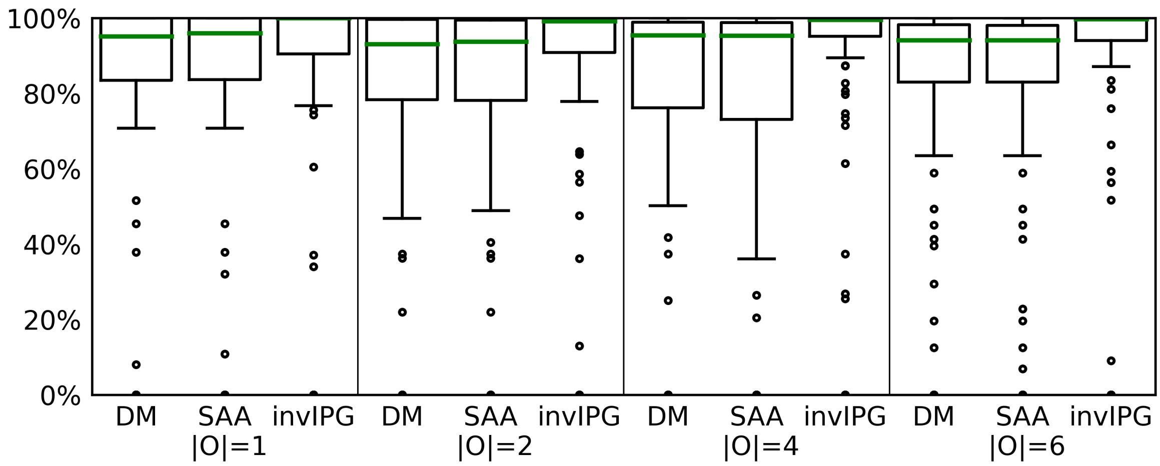

Figure 5.7: Relative profit of a new entrant based on parameters estimated by the invIPG, Distribution mean (DM), or Sample average approximation (SAA) approach for $|J| = |K| = 10$, $n = 100$ samples per boxplot.

The invIPG outperforms both the naïve DM approach and the SAA, with a median profit close to or equal to the profit under full information (100%). Across these instances, using the invIPG leads to an increase in median profits by 4-11 percentage points. The advantage of the invIPG is on the lower end of this range (≈ 4%) for small problem instances ($|J| = |K| = 10$, Figure 5.7) and is particularly evident (≈ 11%) with a larger problem size ($|J| = |K| = 20$). The larger problem size increases the importance of the location decision and thus emphasizes the value of the estimation of customer choice parameters through the invIPG. In addition to this general improvement in the median relative profit, the invIPG clearly benefits from an increase in the number of observations, reducing the occurrence of outliers and improving the reliability of the estimation. This effect is not as apparent for the larger problem size in Figure 5.8 as the invIPG exceeds the 20k seconds timelimit for approximately 40% of all instances and is thus not solved to optimality in many cases with a large amount of observations.

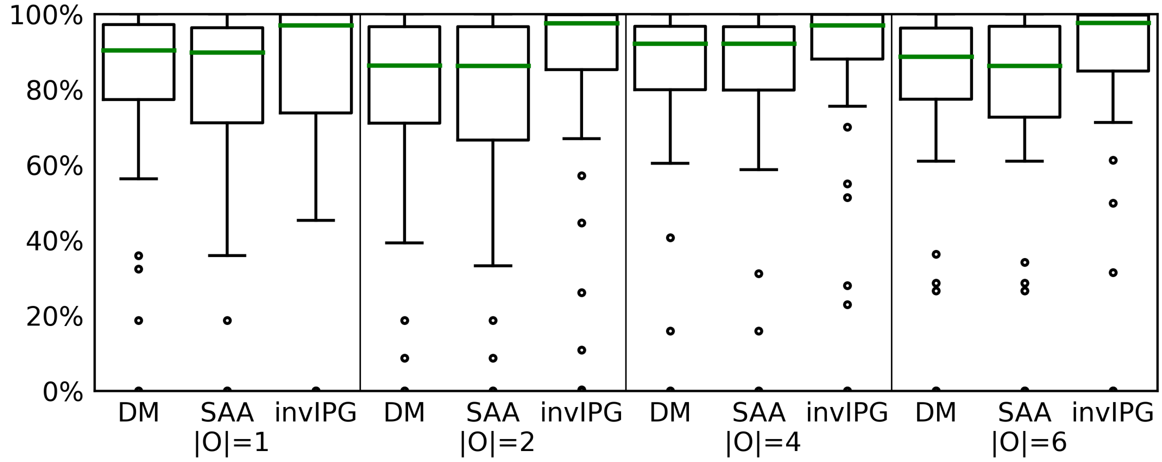

Figure 5.8: Relative profit of a new entrant based on parameters estimated by the invIPG, Distribution mean (DM), or Sample average approximation (SAA) approach for $|J| = |K| = 20$, $n = 40$ samples per boxplot.

5.5 Conclusion

Applying the invIPG method to the competitive retail location problem enables retrieving reliable and helpful information on customer choice. This information particularly benefits new market entrants who can improve their market entry location strategy accordingly. In our numerical experiments, new entrants who observe prior equilibria between incumbents can consistently realize savings of $4 - 11\%$ on average with some instances showing significantly higher improvements compared to a scenario-based approach (SAA). The retrieved customer choice information could also be valuable to other market participants without access to granular sales data, e.g., for city planners, policymakers, or regulatory authorities.

The application of the invIPG towards retail location involves some limitations. In our numerical experiments, we assume that all observations belong to similar regions and that (in the absence of noisy observations) customers behave identically in all regions. It would be valuable to test this hypothesis in future research, leveraging real-world data sets to obtain additional insights into the possible gains of inverse optimization for market entrants in practice. Furthermore, the presented approach relies on parameterizing the customers' utility function. It thus presumes a prior general knowledge of the decision-making of the customer and the incumbents. In future work, one could attempt to learn the customers' appreciation of a store (utility) from the incumbents' decisions using nonparametric machine learning models rather than presupposing the customers' utility function structure. Lastly, the approach assumes a full-information game between incumbents and focuses on information extraction through a new entrant. Similarly, incumbents who possess only partial information about their competitors could aim to improve their competitor knowledge through inverse optimization of observed decisions.

We chose the (inverse) retail location problem as a tangible example of competitive integer decision-making in practice. However, our approach straightforwardly extends to other inverse IPGs. Applications to other problems such as capacity or assortment competition in retail or competition between logistics service providers routing vehicles could prove an exciting avenue for future research.

Appendix 5.A Linearization of the forward problem

Most solvers do not support the fractional, non-convex equality constraints in (5.5)-(5.6). We therefore reformulate (5.6) into the non-convex quadratic inequality

$$\Pi_i^o(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o) \leq \sum_{j \in J} m_{ij}^o p_j^o \hat{f}_{ij}^o - \sum_{k \in K} \hat{x}_{ik}^o c_{ik}^o. \tag{5.21}$$

With f_{ij}^o (the market share for $(\mathbf{x}_i^o, \hat{\mathbf{x}}_{-i}^o)$) following from (5.5), we reformulate the above inequalities to

$$f_{ij}^o \sum_{k \in K | d_{jk} \leq \bar{d}} \left(\sum_{\bar{i} \neq i} \left(\hat{x}_{\bar{i}k} u_{\bar{i}jk}^o \right) + x_{ik}^o u_{ijk}^o \right) \leq \sum_{k \in K | d_{jk} \leq \bar{d}} x_{ik}^o u_{ijk}^o \tag{5.22}$$

$$f_{ij}^o \leq M_{ij}^o \sum_{k \in K | d_{jk} \leq \bar{d}} x_{ik}^o u_{ijk}^o. \tag{5.23}$$

$M_{ij}^o = \left(\min_{k \in K} u_{ijk}^o \right)^{-1}$ provides tight bounds for (5.23). We linearize (5.22) by introducing the auxiliary variable $y_{ijk}^o = f_{ij}^o x_{ik}^o$, yielding

$$\sum_{k \in K | d_{jk} \leq \bar{d}} y_{ijk}^o u_{ijk}^o + \sum_{\bar{i} \neq i} \sum_{k \in K | d_{jk} \leq \bar{d}} f_{ij}^o \hat{x}_{\bar{i}k}^o u_{\bar{i}jk}^o \leq \sum_{k \in K | d_{jk} \leq \bar{d}} x_{ik}^o u_{ijk}^o. \tag{5.24}$$

Because y_{ijk}^o is a product of a binary (x_{ik}^o) and a continuous ($f_{ij}^o \in [0,1]$) decision variable, we can equally represent it based on the following linear inequalities

$$y_{ijk}^o \leq x_{ik}^o \tag{5.25}$$

$$y_{ijk}^o \leq f_{ij}^o \tag{5.26}$$

$$y_{ijk}^o \geq f_{ij}^o - (1 - x_{ik}^o). \tag{5.27}$$

Appendix 5.B Linearization of non-convex linking constraints

The non-linearity in (5.22) arises through multiplication of f_{ij}^o with $\alpha, \boldsymbol{\beta}$ (decision variables in the master problem). This non-convex quadratic problem formulation can be solved by some commercial solvers (e.g., Gurobi 9.5) at the expense of long runtimes. Alternatively, since α,$\boldsymbol{\beta}$ are discrete, with a predetermined number of discrete steps

$(|A|, |B|)$, we can linearize (5.22) using binary expansion (see, e.g. Kleinert et al., 2021). We illustrate the approach for the product of f^o_{ij} and β_i. We represent β_i using binary decision variables $\tilde{\beta}_{si}$ as

$$\beta_i = \frac{\sum_{s\in[0,\ldots,\lceil \ln|B|+1\rceil]} \tilde{\beta}_{si} 2^s}{|B|} \qquad \forall i \in I \tag{5.28}$$

$$\tilde{\beta}_{si} \in \{0,1\} \qquad \forall i \in I,\ \forall s \in [0,\ldots,\lceil \ln|B|+1\rceil] \tag{5.29}$$

Hence, the product $f^o_{ij}\beta_i$ reduces to the product between a binary and continuous variable, which can easily be linearized as

$$f^o_{ij}\beta_i = \frac{\sum_{s\in[0,\ldots,\lceil \ln|B|+1\rceil]} f^o_{ij}\tilde{\beta}_{si} 2^s}{|B|}. \tag{5.30}$$

Appendix 5.C Extended flowchart for the usage of approximate cuts

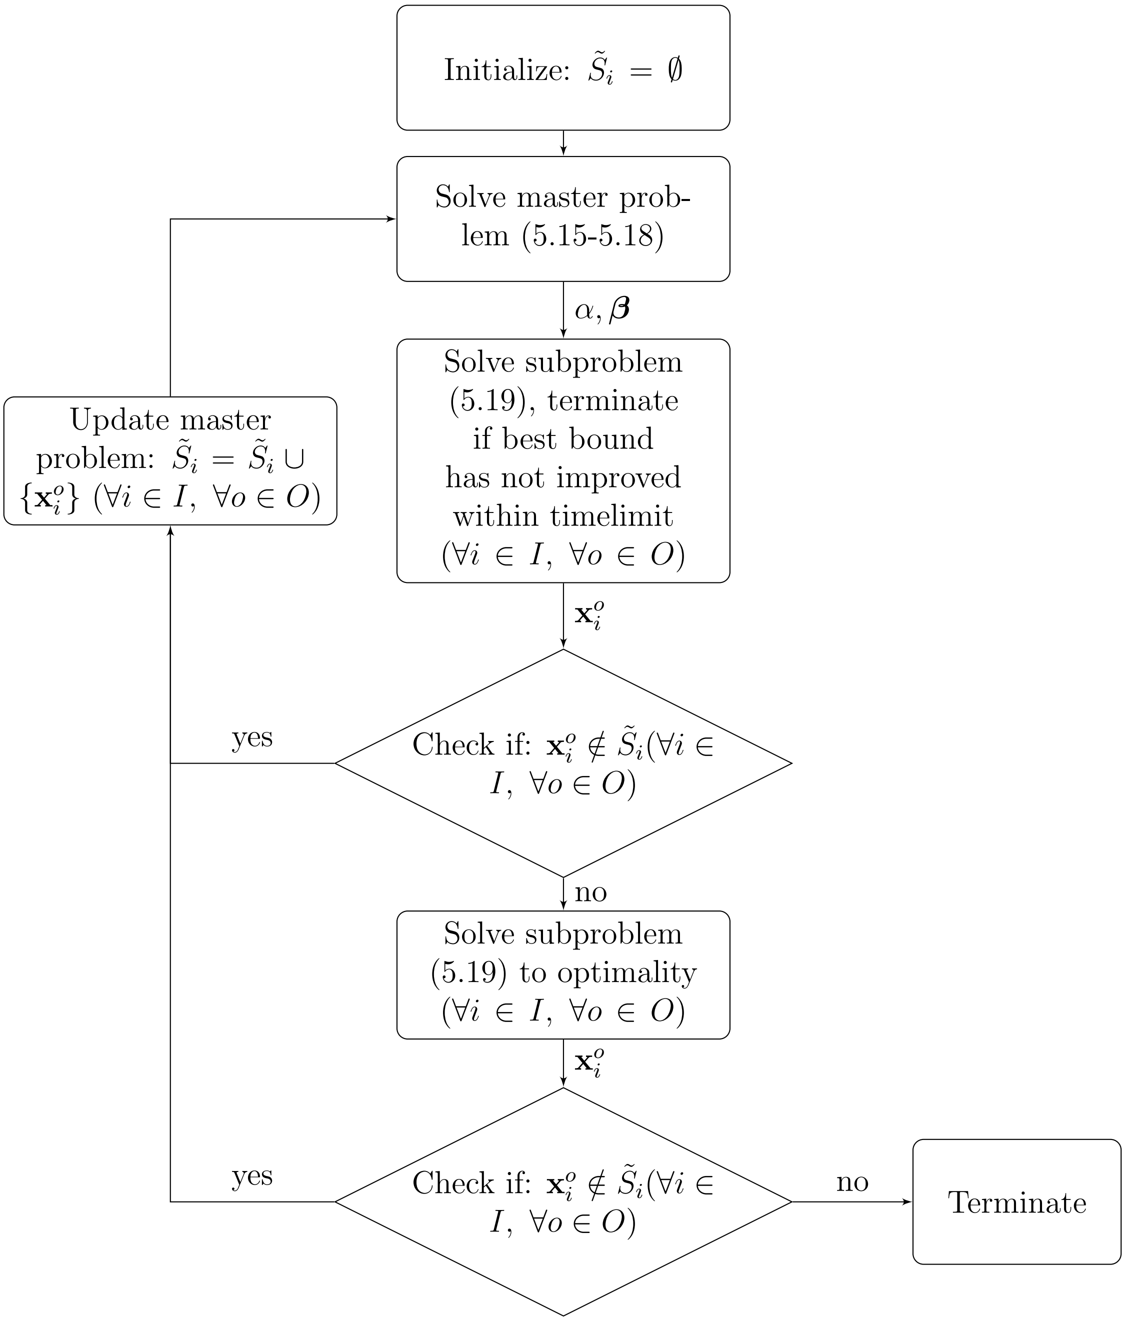

Figure 5.9: Flowchart of invIPG algorithm using approximate cuts

Appendix 5.D Overview of notation

Table 5.2: Sets, indices, parameters, and variables in the inverse problem

Sets

I	set of incumbent retailers
J	set of customer locations (multiple customers per location)
K	set of potential facility locations across all retailers
$K_i \subseteq K$	set of potential facility locations per retailer $i \in I$
O	set of observations
S_i	set of all possible strategies of retailer $i \in I$

Exogenous parameters

p_j^o	population/customer count for $j \in J$ in observation $o \in O$
m_{ij}^o	margin per customer for $j \in J$ for retailer $i \in I$ in observation $o \in O$
c_{ik}^o	costs for retailer $i \in I$ in location $k \in K$ (e.g., rent, salaries, equipment) in observation $o \in O$
$\bar{d}$	maximum distance tolerated by customers
d_{jk}^o	distance between $j \in J$ and location $k \in K$ in observation $o \in O$
$\tilde{d}_{jk}^o$	normalized accessibility between $j \in J$ and location $k \in K$ in observation $o \in O$
g_k^o	convenience of location $k \in K$ based on points of interest in proximity in observation $o \in O$
$\tilde{g}_k^o$	normalized convenience of location $k \in K$ in observation $o \in O$
$\bar{g}$	highest convenience, $\bar{g} = \max_{k \in K, o \in O} g_k^o$
$\hat{\mathbf{x}}_i^o$	observed locations of retailer $i \in I$ in observation $o \in o$
A	set of potential customer choice paramaters $\alpha \in A$
B	set of potential customer choice paramaters $\beta_i \in B$

Endogenous variables

u_{ijk}^o	utility of facility $k \in K$ of $i \in I$ for customers $j \in J$
f_{ij}^o	fraction of customers in j patronizing retailer i
x_{ik}^o	1, only if retailer $i \in I$ opens facility $k \in K$
α	normalized sensitivity towards distance
$\boldsymbol{\beta} = (\beta_i)_{i \in I}$	vector of brand attractiveness for retail chains $i \in I$
$\boldsymbol{\delta} := (\delta_i^o)_{i \in I, o \in O}$	vector of unilateral improvement potentials

Appendix 5.E Hamming distances as an alternative objective function

In Sections 5.2.5 and 5.3.2, we define the minimization of the cumulative unilateral improvement potential ($\min_{\alpha,\boldsymbol{\beta}} ||\boldsymbol{\delta}||$) as the overarching objective of the invIPG. Instead of minimizing this difference in profits between a decision $\tilde{\mathbf{x}}_i^o$ that is optimal for a parameter set $(\alpha, \boldsymbol{\beta})$ and an observed decision ($\hat{\mathbf{x}}_i^o$), one could minimize the Hamming distance (HD) between the two decisions, i.e.,

$$\min_{\alpha,\boldsymbol{\beta}} \left|\left|(\tilde{\mathbf{x}}_i^o - \hat{\mathbf{x}}_i^o)_{i\in I, o\in O}\right|\right| \tag{5.31}$$

$$\text{s.t.} \qquad \Pi_i^o(\tilde{\mathbf{x}}_i^o, \hat{\mathbf{x}}_i^o) \geq \max_{\mathbf{x}_i^o \in S_i} \Pi_i^o(\mathbf{x}_i^o, \hat{\mathbf{x}}_i^o) \qquad \forall i \in I, \forall o \in O. \tag{5.32}$$

In Figure 5.10, we refer to this alternative as the invIPG-HD and compare runtime and solution quality with the invIPG. As a metric for the solution quality, we use the error (L_2-norm) between the inversely estimated parameters and the ground-truth values (sampled from distributions as in Section 5.4.3). While solution quality is almost identical for both objective functions, the invIPG based on the unilateral improvement potential presented in Section 5.2.5 and 5.3.2 results in approximately 2-times faster average runtimes compared with the invIPG-HD. This computational advantage of the invIPG can be explained by the number of binary decision variables in the master problem. The invIPG represents a strict decomposition; all (binary) location decisions are outsourced to the subproblems. The master problem merely compares payoffs for enumerated subproblem solutions with the observed location structure. In contrast, the invIPG-HD requires some binary decisions decision variables $\tilde{\mathbf{x}}_i^o$ in the master problem, significantly increasing its complexity and the number of nodes the solver explores in its branch and bound process.

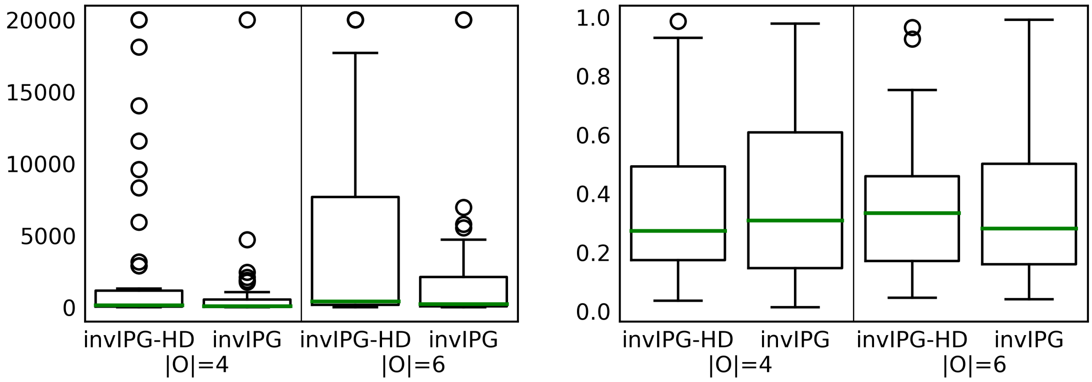

(a) Runtime in seconds

(b) Error (L_2-norm) between inversely estimated parameter and ground-truth values

Figure 5.10: Effects of a Hamming distance based objective function (invIPG-HD) on runtime and solution quality for problem size $|J| = |K| = 10$ and $n = 50$ samples per boxplot

Chapter 6

Conclusion

6.1 Summary

We investigate different aspects of simultaneous competitive decision-making in finite games and extend prior methods to enable the identification, selection, and inverse optimization of equilibria in this class of games. In various numerical experiments across different applications, we quantify the advantage of acting in anticipation of competition.

In Chapter 3, we propose the exhaustive Sample generation method (eSGM) for the identification of all equilibria and the subsequent equilibrium selection in finite games. The eSGM builds on the existing Sample generation method (SGM) that enables the identification of a single equilibrium.

RQ 1.1 *How can we identify all equilibria in a finite game?*

The eSGM complements the sampling approach of the SGM with two key changes. Instead of identifying a single equilibrium in the current sample, we identify all equilibria using an adaptation of MIP-Nash for n-player normal-form games. Rather than only enlarging the sample by best-responses to these equilibria, we also add further candidate solutions based on necessary but not sufficient criteria for additional equilibria. Jointly, these two changes enable the identification of all equilibria.

RQ 1.2 *Having identified all equilibria in a finite game – how can we select the equilibrium that is expected to be the outcome of the game based on the respective player incentives and the equilibrium selection theory by Harsanyi, 1995?*

We show that, after the enumeration of all equilibria through the eSGM, the equilibrium selection theory by Harsanyi can be applied with minor adaptations. As the approach requires complex high-dimensional volume calculations of non-linearly bounded sets, we propose a Monte Carlo volume estimation as an alternative to exact volume calculations. In various numerical examples of knapsack, location and lot-sizing games, we show the advantage of the eSGM with equilibrium selection over its predecessor (SGM). For example, in the examined instances of a 2-player knapsack game the equilibrium selected through the eSGM on average has a 40% higher probability compared with the equilibrium identified through the SGM.

In Chapter 4, we apply the eSGM to a case study of Hydrogen fuel station (HFS) location under emerging competition.

RQ 2.1 *How will competition between station providers influence the emerging hydrogen refueling network structure?*

We extend prior model formulations for alternative fuel station location selection to this simultaneous competitive setting. Depending on the network structure and economic parameters, one can observe concentration effects, as reported by Hotelling, 1929 for competitive location selection on a line, or dispersion effects.

RQ 2.2 *How valuable is it for decision makers to take competitor actions into consideration?*

In a case study for Munich, we show that decision makers acting in anticipation of apparent competition can increase their profits by 17% on average when compared with decision makers neglecting competition.

RQ 2.3 *Should policymakers foster (e.g., through government-backed provider associations) or impede (e.g., through strict antitrust laws) collaboration between competing providers?*

Collaboration between the competitors could further increase profits by 28% on average, however, the sparser network of HFSs would imply longer deviation distances for customers with detours increasing by 3% on average.

Last, in Chapter 5, we devise a method for the inverse optimization of parameters that lead to observed equilibria in Integer programming games (IPGs).

RQ 3.1 *How can we identify parameters that lead to observed equilibria in situations of simultaneous competition using inverse optimization, taking into account integrality constraints?*

We reformulate the inverse optimization problem for IPGs into a bilevel problem. Here, the master problem selects parameters that minimize the unilateral improvement potential for all players. To determine this unilateral improvement potential, we compare the players' payoffs in the observed equilibrium with their payoffs in alternative strategies, which are enumerated iteratively in subproblems.

RQ 3.2 *How valuable are the derived parameter estimates to new entrants, optimizing their own market entry location strategy?*

The advantage of this approach becomes particularly evident in an exemplary application where a new entrant lacks detailed information on underlying problem parameters, but has knowledge of the general payoff structure of incumbents and can observe the equilibrium outcome of their competitive location selection. For the competitive retail location problem, we show that a new entrant using our inverse optimization approach to extract hidden information on customer choice parameters can improve their profits by 4-11%, when compared with a new entrant relying on simple statistical information or sample average approximation.

6.2 Limitations and future research

The methods discussed in this thesis focus on finite Mathematical programmming games (MPGs) or finite IPGs. By definition, a mixed-integer game in which some decisions are continuous is not finite and cannot be addressed by our methodology. In the absence of a finite set of strategies for each decision maker, the original proof of existence of Nash equilibria (Nash, 1951) does not hold. While there are some conditions under which an equilibrium is guaranteed to exist for non-finite games (see, e.g., Carvalho et al., 2022), some of the theorems presented throughout this thesis would not hold. The adaptation of methods presented in this thesis to non-finite MPGs thus provides a valuable future research opportunity.

Whereas the applied equilibrium selection approach based on the theory of Harsanyi (1995) does intuitively hold for rational, non-cooperative decision makers, decision making in practice can be far from rational. A detailed comparison of behavioral analyses of competitive decision making (see e.g., Camerer, 2011) and predicted outcomes based on the presented methodology would be required to confirm model outcomes in practice.

We apply the methods toward competitive location selection for alternative fuel stations and retailing. In both cases, we assume a simultaneous one-shot game between the competitors, whereas in practice the location decisions will likely be dispersed over several years. An extension of both solving methodology and model formulations to multi-stage games could thus further improve practical applicability.

Such a dynamic multi-stage game could also include uncertainties beyond the competitors' decision making. For example, in the case of HFS location selection in Chapter 4, after their initial location selection in the first stage, providers might gain additional knowledge on Fuel cell electric vehicle (FCEV) penetration and customer demand to use in a potential second stage. It might thus become necessary to extend the deterministic perspective discussed in this work with stochastic considerations.

The applications and model formulations addressed in this thesis are of considerable size, and strategy sets can encompass several billion possible strategies per player. However, at this problem size, the proposed methods reach computational limitations with runtimes approaching multiple hours and in rare cases exceeding a full day. Hence, larger-scale problems will likely require additional algorithmic advances or (potentially problem-specific) decomposition approaches to enable reasonable runtimes.

While both the eSGM proposed in Chapter 3 and the invIPG discussed in Chapter 5 are formulated for n-player games, for the eSGM practical application is limited to a small number of players. Already games with $n = 3$ players require a solver for quadratic constraints, games with $n > 3$ players would need reformulation and will be hard to solve.

Despite the above limitations, the methods discussed in this thesis could already yield valuable theoretical and managerial insights in other applications of integer/binary problems. The majority of these problems are as of yet considered without addressing competition, or competitive considerations are limited to leader-follower relationships or continuous relaxations of these problems. Among others, examples of potential further applications include inventory problems, capacity decisions, or assortment competition.

Bibliography

Aboolian, Robert, Oded Berman, and Dmitry Krass (2007). Competitive facility location and design problem. *European Journal of Operational Research* 182 (1), pp. 40–62.

Access Development (2016). *The Impact of Retail Proximity on Consumer Purchases*. Tech. rep. Last accessed 17.03.2022. URL: `https://cdn2.hubspot.net/hubfs/263750/Access_Consumer_Spend_Study_2016.pdf`.

Ahuja, Ravindra K and James B Orlin (2001). Inverse optimization. *Operations Research* 49 (5), pp. 771–783.

Allen, Stephanie, Steven A Gabriel, and John P Dickerson (2022). Using inverse optimization to learn cost functions in generalized Nash games. *Computers & Operations Research* 142. Article 105721.

Aloise, Daniel, Amit Deshpande, Pierre Hansen, and Preyas Popat (2009). NP-hardness of Euclidean sum-of-squares clustering. *Machine Learning* 75 (2), pp. 245–248.

Anders, Sascha (2015). Discounter and supermarket: Research about impacts on traffic, catchment areas, customer preferences and approval processes against the background of § 11 (3) federal land utilisation ordinance. *Raumforschung und Raumordnung — Spatial Research and Planning* 73 (3), pp. 219–232.

Anderson, Edward, Bo Chen, and Lusheng Shao (2017). Supplier competition with option contracts for discrete blocks of capacity. *Operations Research* 65 (4), pp. 952–967.

Apostolou, Dimitrios and George Xydis (2019). A literature review on hydrogen refuelling stations and infrastructure. Current status and future prospects. *Renewable and Sustainable Energy Reviews* 113. Article 109292.

Arslan, Okan, Oya E Karaşan, Ali R Mahjoub, and Hande Yaman (2019). A branch-and-cut algorithm for the alternative fuel refueling station location problem with routing. *Transportation Science* 53 (4), pp. 1107–1125.

Arthur, David and Sergei Vassilvitskii (2006). *k-means++: The advantages of careful seeding*. Tech. rep. Last accessed 12.07.2022. Stanford. URL: `http://ilpubs.stanford.edu:8090/778/1/2006-13.pdf`.

Aswani, Anil, Zuo-Jun Shen, and Auyon Siddiq (2018). Inverse optimization with noisy data. *Operations Research* 66 (3), pp. 870–892.

— (2019). Data-driven incentive design in the medicare shared savings program. *Operations Research* 67 (4), pp. 1002–1026.

Aumann, Robert J (1990). Nash equilibria are not enforceable. *Economic Decision-Making: Games, Econometrics, and Optimization.* Ed. by JJ Gabszewicz, J.-F. Richard, and LA Wolsey. Amsterdam: Elsevier, pp. 201–206.

Bärmann, Andreas, Sebastian Pokutta, and Oskar Schneider (2017). Emulating the expert: Inverse optimization through online learning. *ICML'17: Proceedings of the 34th International Conference on Machine Learning* 70, pp. 400–410.

Barnhart, Cynthia, Ellis L Johnson, George L Nemhauser, Martin W P Savelsbergh, and Pamela H Vance (1998). Branch-and-price: Column generation for solving huge integer programs. *Operations Research* 46 (3), pp. 293–432.

Belavina, Elena, Karan Girotra, and Ashish Kabra (2017). Online grocery retail: Revenue models and environmental impact. *Management Science* 63 (6), pp. 1781–1799.

Bell, David R, Teck-Hua Ho, and Christopher S Tang (1998). Determining where to shop: Fixed and variable costs of shopping. *Journal of Marketing Research* 35 (3), pp. 352–369.

Ben-Akiva, Moshe and Steven R Lerman (1997). *Discrete choice analysis*. 7th. Cambridge, Massachusetts: The MIT Press.

Benati, Stefano and Pierre Hansen (2002). The maximum capture problem with random utilities: Problem formulation and algorithms. *European Journal of Operational Research* 143 (3), pp. 518–530.

Berbeglia, Gerardo, Agustín Garassino, and Gustavo Vulcano (2021). A comparative empirical study of discrete choice models in retail operations. *Management Science* 68 (6), pp. 4005–4023.

Beresnev, V. (n.d.). Competitive facility location and design problem. Last accessed 11.07.2020. URL: `http://www.math.nsc.ru/AP/benchmarks/Design/design_en.html`.

Bergantino, Angela S, Claudia Capozza, and Mario Intini (2020). Empirical investigation of retail fuel pricing: The impact of spatial interaction, competition and territorial factors. *Energy Economics* 90. Article 104876.

Berman, Oded, Dimitris Bertsimas, and Richard C Larson (1995a). Locating discretionary service facilities, II: Maximizing market size, minimizing inconvenience. *Operations Research* 43 (4), pp. 623–632.

Berman, Oded and Dmitry Krass (1998). Flow intercepting spatial interaction model: A new approach to optimal location of competitive facilities. *Location Science* 6 (1-4), pp. 41–65.

Berman, Oded, Dmitry Krass, and Chen W Xu (1995b). Locating discretionary service facilities based on probabilistic customer flows. *Transportation Science* 29 (3), pp. 276–290.

Bersani, Chiara, Riccardo Minciardi, Roberto Sacile, and Eva Trasforini (2009). Network planning of fuelling service stations in a near-term competitive scenario of the hydrogen economy. *Socio-Economic Planning Sciences* 43 (1), pp. 55–71.

Bertsimas, Dimitris, Vishal Gupta, and Ioannis C Paschalidis (2015). Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming* 153 (2), pp. 595–633.

Birge, John R and François Louveaux (2011). *Introduction to stochastic programming.* Ed. by Thomas V. Mikosch, Sidney I. Resnick, and Stephen M. Robinson. 2nd. Springer Series in Operations Research and Financial Engineering. New York, NY: Springer New York.

Bodur, Merve, Timothy CY Chan, and Ian Yihang Zhu (2022). Inverse mixed integer optimization: Polyhedral insights and trust region methods. *INFORMS Journal on Computing* 34 (3), pp. 1471–1488.

Brown, Tim, Lori Smith, Shane Stephens-Romero, and Scott Samuelsen (2013). Economic analysis of near-term California hydrogen infrastructure. *International Journal of Hydrogen Energy* 38 (10), pp. 3846–3857.

Cachon, Gérard P and Paul H Zipkin (1999). Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science* 45 (7), pp. 936–953.

Camerer, Colin F (2011). *Behavioral game theory: Experiments in strategic interaction.* Princeton University Press.

Capar, Ismail and Michael Kuby (2012). An efficient formulation of the flow refueling location model for alternative-fuel stations. *IIE Transactions (Institute of Industrial Engineers)* 44 (8), pp. 622–636.

Carvalho, Margarida, Gabriele Dragotto, Andrea Lodi, and Sriram Sankaranarayanan (2021). The cut and play algorithm: Computing Nash equilibria via outer approximations. *arXiv preprint* arXiv:2111.05726.

Carvalho, Margarida, Andrea Lodi, and João P Pedroso (2018a). Existence of Nash equilibria on integer programming games. *Springer Proceedings in Mathematics and Statistics* 223, pp. 11–23.

Carvalho, Margarida, Andrea Lodi, and Joao P Pedroso (2022). Computing Nash equilibria for integer programming games. *European Journal of Operational Research* 303 (3), pp. 1057–1070.

Carvalho, Margarida, João P Pedroso, Claudio Telha, and Mathieu Van Vyve (2018b). Competitive uncapacitated lot-sizing game. *International Journal of Production Economics* 204, pp. 148–159.

Chan, Timothy CY, Rafid Mahmood, and Ian Y Zhu (2021). Inverse optimization: Theory and applications. *arXiv preprint* arXiv:2109.03920.

Charness, Gary, Francesco Feri, Miguel A Meléndez-Jiménez, and Matthias Sutter (2014). Experimental games on networks: Underpinnings of behavior and equilibrium selection. *Econometrica* 82 (5), pp. 1615–1670.

Chen, Lu, Yuyi Chen, and André Langevin (2021). An inverse optimization approach for a capacitated vehicle routing problem. *European Journal of Operational Research* 295 (3), pp. 1087–1098.

Chung, Sung H and Changhyun Kwon (2015). Multi-period planning for electric car charging station locations: A case of Korean expressways. *European Journal of Operational Research* 242 (2), pp. 677–687.

Cleeren, Kathleen, Frank Verboven, Marnik G Dekimpe, and Katrijn Gielens (2010). Intra- and interformat competition among discounters and supermarkets. *Marketing Science* 29 (3), pp. 456–473.

Cooper, Lee G and Masao Nakanishi (1988). *Market-share analysis: Evaluating competitive marketing effectiveness.* Boston, Dorbrecht, London: Kluwer academic publishers.

Cooper, Leon (1963). Location-allocation problems. *Operations Research* 11 (3), pp. 331–343.

Crönert, Tobias, Layla Martin, Stefan Minner, and Christopher S. Tang (2022). Inverse optimization of integer programming games for parameter estimation arising from competitive retail location selection. *Available at SSRN* 4147765.

Crönert, Tobias and Stefan Minner (2021a). Equilibrium identification and selection in finite games. *Available at SSRN* 3762380.

— (2021b). Location selection for hydrogen fuel stations under emerging provider competition. *Transportation Research Part C: Emerging Technologies* 133. Article 103426.

Dagdougui, Hanane, Rorberto Sacile, Chiara Bersani, and Ahmed Ouammi (2018). Hydrogen logistics: Safety and risks issues. *Hydrogen Infrastructure for Energy Applications.* London: Academic Press, pp. 127–148.

Daskalakis, Constantinos, Aranyak Mehta, and Christos Papadimitriou (2006). A note on approximate Nash equilibria. *International Workshop on Internet and Network Economics*. Springer, pp. 297–306.

Dixit, Avinash (1980). The role of investment in entry-deterrence. *The Economic Journal* 90 (357), pp. 95–106.

Dobson, Gregory and Uday S Karmarkar (1987). Competitive location on a network. *Operations Research* 35 (4), pp. 565–574.

Dragotto, Gabriele, Sriram Sankaranarayanan, Margarida Carvalho, and Andrea Lodi (2021). ZERO: Playing mathematical programming games. *arXiv preprint* arXiv:2111.07932.

Drezner, Tammy (2014). A review of competitive facility location in the plane. *Logistics Research* 7. Article 114.

Drezner, Tammy, Zvi Drezner, and Pawel Kalczynski (2020). Gradual cover competitive facility location. *OR Spectrum* 42 (2), pp. 333–354.

Dyer, Martin E and Alan M Frieze (1988). On the complexity of computing the volume of a polyhedron. *SIAM Journal on Computing* 17 (5), pp. 967–974.

Egri, Péter, Tamás Kis, András Kovács, and József Váncza (2014). An inverse economic lot-sizing approach to eliciting supplier cost parameters. *International Journal of Production Economics* 149, pp. 80–88.

Eiselt, Horst A and Gilbert Laporte (1997). Sequential location problems. *European Journal of Operational Research* 96 (2), pp. 217–231.

Eiselt, Horst A, Vladimir Marianov, and Tammy Drezner (2019). Competitive location models. *Location Science*. Ed. by Gilbert Laporte, Stefan Nickel, and Francisco Saldanha da Gama. Springer, pp. 391–429.

European Commission (2020). *A hydrogen strategy for a climate-neutral Europe*. Tech. rep. Last accessed 12.07.2022. URL: `https://ec.europa.eu/energy/sites/ener/files/hydrogen_strategy.pdf`.

European Environment Agency (2020). *EEA greenhouse gas*. Tech. rep. Last accessed 12.07.2022. URL: `https://www.eea.europa.eu/data-and-maps/data/data-viewers/greenhouse-gases-viewer`.

Facchinei, Francisco and Christian Kanzow (2007). Generalized Nash equilibrium problems. *4OR* 5 (3), pp. 173–210.

Federgruen, Awi and Ming Hu (2015). Multi-product price and assortment competition. *Operations Research* 63 (3), pp. 572–584.

Feldman, Michal and Tami Tamir (2012). Conflicting congestion effects in resource allocation games. *Operations Research* 60 (3), pp. 529–540.

Fok, Danny SK and Daniel Crevier (1989). Volume estimation by monte carlo methods. *Journal of Statistical Computation and Simulation* 31 (4), pp. 223–235.

Friedman, Daniel and Shuchen Zhao (2021). When are mixed equilibria relevant? *Journal of Economic Behavior & Organization* 191, pp. 51–65.

Friesz, Terry L and David Bernstein (2016). Nash games. *Complex Networks and Dynamic Systems 3: Foundations of Network Optimization and Games.* Boston: Springer, pp. 265–323.

Gilboa, Itzhak and Eitan Zemel (1989). Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior* 1 (1), pp. 80–93.

Godinho, Pedro and Joana Dias (2010). A two-player competitive discrete location model with simultaneous decisions. *European Journal of Operational Research* 207 (3), pp. 1419–1432.

Guo, Zhaomiao, Julio Deride, and Yueyue Fan (2016). Infrastructure planning for fast charging stations in a competitive market. *Transportation Research Part C: Emerging Technologies* 68, pp. 215–227.

Haase, Knut and Sven Müller (2013). Management of school locations allowing for free school choice. *Omega* 41 (5), pp. 847–855.

Harsanyi, John C (1995). A new theory of equilibrium selection for games with complete information. *Games and Economic Behavior* 10 (2), pp. 91–122.

Harsanyi, John C and Reinhard Selten (1988). *A general theory of equilibrium selection in games.* Cambridge, Massachusetts: The MIT Press.

He, Jia, Hai Yang, Tie Qiao Tang, and Hai Jun Huang (2018). An optimal charging station location model with the consideration of electric vehicle's driving range. *Transportation Research Part C: Emerging Technologies* 86, pp. 641–654.

Hemmecke, Raymond, Shmuel Onn, and Robert Weismantel (2009). Nash-equilibria and N-fold integer programming. *arXiv preprint* arXiv:0903.4577.

Henk, Martin, Jürgen Richter-Gebert, and Günter M Ziegler (2017). Basic properties of convex polytopes. *Handbook of Discrete and Computational Geometry.* Ed. by Jacob E. Gooodman, Joseph O'Rourke, and Csaba D. Toth. 3rd. Boca Raton: Chapman and Hall/CRC, pp. 243–270.

Heuberger, Clemens (2004). Inverse combinatorial optimization: A survey on problems, methods, and results. *Journal of Combinatorial Optimization* 8 (3), pp. 329–361.

Hodgson, John M (1990). A flow capturing location allocation model. *Geographical Analysis* 22 (3), pp. 270–279.

Hosseini, Meysam and Seyyed A MirHassani (2015). Refueling-station location problem under uncertainty. *Transportation Research Part E: Logistics and Transportation Review* 84, pp. 101–116.

Hotelling, Harold (1929). Stability in competition. *The Economic Journal* 39 (153), pp. 41–57.

Huang, Yongxi, Shengyin Li, and Zhen Sean Qian (2015). Optimal deployment of alternative fueling stations on transportation networks considering deviation paths. *Networks and Spatial Economics* 15 (1), pp. 183–204.

Huff, David L (1964). Defining and estimating a trade area. *Journal of Marketing* 28 (3), pp. 34–38.

Huppmann, Daniel and Sauleh Siddiqui (2018). An exact solution method for binary equilibrium problems with compensation and the power market uplift problem. *European Journal of Operational Research* 266 (2), pp. 622–638.

Jung, Jaeyoung, Joseph Y.J. Chow, R. Jayakrishnan, and Ji Young Park (2014). Stochastic dynamic itinerary interception refueling location problem with queue delay for electric taxi charging stations. *Transportation Research Part C: Emerging Technologies* 40, pp. 123–142.

Kim, Jong-Geun and Michael Kuby (2011). The deviation-flow refueling location model for optimizing a network of refueling stations. *International Journal of Hydrogen Energy* 37 (6), pp. 5406–5420.

Kleinert, Thomas, Veronika Grimm, and Martin Schmidt (2021). Outer approximation for global optimization of mixed-integer quadratic bilevel problems. *Mathematical Programming* 188, pp. 461–521.

Ko, Joonho, Tae Hyoung Tommy Gim, and Randall Guensler (2017). Locating refuelling stations for alternative fuel vehicles: A review on models and applications. *Transport Reviews* 37 (5), pp. 551–570.

Köhler, Jonathan, Martin Wietschel, Lorraine Whitmarsh, Dogan Keles, and Wolfgang Schade (2010). Infrastructure investment for a transition to hydrogen automobiles. *Technological Forecasting and Social Change* 77 (8), pp. 1237–1248.

Köppe, Matthias, Christopher Thomas Ryan, and Maurice Queyranne (2011). Rational generating functions and integer programming games. *Operations Research* 59 (6), pp. 1445–1460.

Koutsoupias, Elias and Christos Papadimitriou (1999). Worst-case equilibria. *STACS 99*. Ed. by Christoph Meinel and Sophie Tison. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 404–413.

Kuby, Michael and Seow Lim (2005). The flow-refueling location problem for alternative-fuel vehicles. *Socio-Economic Planning Sciences* 39 (2), pp. 125–145.

Kuby, Michael, Lee Lines, Ronald Schultz, Zhixiao Xie, Jong-geun Kim, and Seow Lim (2009). Optimization of hydrogen stations in Florida using the Flow-Refueling Location Model. *International Journal of Hydrogen Energy* 34 (15), pp. 6045–6064.

Kuby, Michael J., Scott B. Kelley, and Joseph Schoenemann (2013). Spatial refueling patterns of alternative-fuel and gasoline vehicle drivers in Los Angeles. *Transportation Research Part D: Transport and Environment* 25, pp. 84–92.

Lemke, Carlton E and Joseph T Howson, Jr. (1964). Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics* 12 (2), pp. 413–423.

Leszczyc, Peter TL Popkowski, Ashish Sinha, and Anna Sahgal (2004). The effect of multi-purpose shopping on pricing and location strategy for grocery stores. *Journal of Retailing* 80 (2), pp. 85–99.

Li, Shengyin, Yongxi Huang, and Scott J. Mason (2016). A multi-period optimization model for the deployment of public electric vehicle charging stations on network. *Transportation Research Part C: Emerging Technologies* 65, pp. 128–143.

Li, Yingru and Lin Liu (2012). Assessing the impact of retail location on store performance: A comparison of Wal-Mart and Kmart stores in Cincinnati. *Applied Geography* 32 (2), pp. 591–600.

Li, Yushan, Fengming Cui, and Lefei Li (2018). An integrated optimization model for the location of hydrogen refueling stations. *International Journal of Hydrogen Energy* 43 (42), pp. 19636–19649.

Lieberman, Marvin B and David B Montgomery (1988). First-mover advantages. *Strategic Management Journal* 9 (S1), pp. 41–58.

Lim, Seow and Michael Kuby (2010). Heuristic algorithms for siting alternative-fuel stations using the Flow-Refueling Location Model. *European Journal of Operational Research* 204 (1), pp. 51–61.

Lin, Cheng Chang and Chuan Chih Lin (2018). The p-center flow-refueling facility location problem. *Transportation Research Part B: Methodological* 118, pp. 124–142.

Lippman, Steven A and Kevin F McCardle (1997). The competitive newsboy. *Operations Research* 45 (1), pp. 54–65.

Lloyd, Stuart P (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28 (2), pp. 129–137.

Luce, Robert D (1959). *Individual choice behavior.* New York: John Wiley.

Lyons, Glenn and Cody Davidson (2016). Guidance for transport planning and policy-making in the face of an uncertain future. *Transportation Research Part A: Policy and Practice* 88, pp. 104–116.

Mahajan, Meena, Prajakta Nimbhorkar, and Kasturi Varadarajan (2012). The planar k-means problem is NP-hard. *Theoretical Computer Science* 442, pp. 13–21.

Marianov, Vladimir, Horst A Eiselt, and Armin Lüer-Villagra (2018). Effects of multi-purpose shopping trips on retail store location in a duopoly. *European Journal of Operational Research* 269 (2), pp. 782–792.

Marianov, Vladimir, Miguel Ríos, and Manuel José Icaza (2008). Facility location for market capture when users rank facilities by shorter travel and waiting times. *European Journal of Operational Research* 191 (1), pp. 32–44.

McFadden, Daniel (1974). Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics.* Ed. by Paul Zarembka. New York: Academic Press, pp. 105–142.

McKinnon, Alan, Michael Browne, Maja Piecyk, and Anthony Whiteing, eds. (2015). *Green logistics.* 3rd. London: Kogan Page.

Milchtaich, Igal (2009). Weighted congestion games with separable preferences. *Games and Economic Behavior* 67 (2), pp. 750–757.

MirHassani, Seyyed A and Roozbeh Ebrazi (2013). A flexible reformulation of the refueling station location problem. *Transportation Science* 47 (4), pp. 617–628.

Moeckel, Rolf, Nico Kuehnel, Carlos Llorca, Ana Tsui Moreno, and Hema Rayaprolu (2020). Agent-based simulation to improve policy sensitivity of trip-based models. *Journal of Advanced Transportation* 2020. Article 1902162.

Moghaddass, Mahsa and Daria Terekhov (2020). Inverse integer optimization with an imperfect observation. *Operations Research Letters* 48 (6), pp. 763–769.

— (2021). Inverse integer optimization with multiple observations. *Optimization Letters* 15 (4), pp. 1061–1079.

Muratori, Matteo (2018). Impact of uncoordinated plug-in electric vehicle charging on residential power demand. *Nature Energy* 3 (3), pp. 193–201.

Muter, Ibrahim, Ilker Birbil, and Kerem Bülbül (2013). Simultaneous column-and-row generation for large-scale linear programs with column-dependent-rows. *Mathematical Programming* 142 (1-2), pp. 47–82.

Myerson, Roger B (1978). Refinements of the Nash equilibrium concept. *International Journal of Game Theory* 7 (2), pp. 73–80.

Nash, John (1951). Non-cooperative games. *The Annals of Mathematics* 54 (2), pp. 286–295.

Netessine, Serguei and Robert A Shumsky (2005). Revenue management games: Horizontal and vertical competition. *Management Science* 51 (5), pp. 813–831.

Nisan, Noam, Tim Roughgarden, Eva Tardos, and Vijay V. Varzirani, eds. (2007). *Algorithmic game theory*. Cambridge University Press.

OSM (2020). Fuel stations in Munich. Last accessed 1.03.2020. URL: `https://www.openstreetmap.org`.

Östling, Robert, Joseph Tao-yi Wang, Eileen Y Chou, and Colin F Camerer (2011). Testing game theory in the field: Swedish LUPI lottery games. *American Economic Journal: Microeconomics* 3 (3), pp. 1–33.

Pagany, Raphaela, Luis Ramirez Camargo, and Wolfgang Dorner (2019). A review of spatial localization methodologies for the electric vehicle charging infrastructure. *International Journal of Sustainable Transportation* 13 (6), pp. 433–449.

Pancras, Joseph, Srinivasaraghavan Sriram, and Vineet Kumar (2012). Empirical investigation of retail expansion and cannibalization in a dynamic environment. *Management Science* 58 (11), pp. 2001–2018.

Pang, Jong-Shi (2010). Three modeling paradigms in mathematical programming. *Mathematical Programming* 125 (2), pp. 297–323.

Papadimitriou, Christos H (1994). On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences* 48 (3), pp. 498–532.

— (2007). The complexity of finding Nash equilibria. *Algorithmic Game Theory*. Ed. by Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Varzirani. Cambridge University Press, pp. 29–51.

Pibernik, Richard, Yingying Zhang, Florian Kerschbaum, and Axel Schröpfer (2011). Secure collaborative supply chain planning and inverse optimization – The JELS model. *European Journal of Operational Research* 208 (1), pp. 75–85.

Plastria, Frank (2001). Static competitive facility location: An overview of optimisation approaches. *European Journal of Operational Research* 129 (3), pp. 461–470.

Porter, Ryan, Eugene Nudelman, and Yoav Shoham (2004). Simple search methods for finding a Nash equilibrium. *Proceedings of the National Conference on Artificial Intelligence*, pp. 664–669.

Ramea, Kalai (2019). An integrated quantitative-qualitative study to monitor the utilization and assess the perception of hydrogen fueling stations. *International Journal of Hydrogen Energy* 44 (33), pp. 18225–18239.

Revelle, Charles S and Horst A Eiselt (2005). Location analysis: A synthesis and survey. *European Journal of Operational Research* 165 (1), pp. 1–19.

Röller, Lars-Hendrik and Mihkel M. Tombak (1993). Competition and investment in flexible technologies. *Management Science* 39 (1), pp. 107–114.

Rosenthal, Robert W (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2 (1), pp. 65–67.

Roughgarden, Tim, ed. (2020). *Beyond the worst-case analysis of algorithms.* Cambridge University Press.

Sagratella, Simone (2016). Computing all solutions of Nash equilibrium problems with discrete strategy sets. *SIAM Journal on Optimization* 26 (4), pp. 2190–2218.

— (2017a). Algorithms for generalized potential games with mixed-integer variables. *Computational Optimization and Applications* 68 (3), pp. 689–717.

— (2017b). Computing equilibria of Cournot oligopoly models with mixed-integer quantities. *Mathematical Methods of Operations Research* 86 (3), pp. 549–565.

— (2019). On generalized Nash equilibrium problems with linear coupling constraints and mixed-integer variables. *Optimization* 68 (1), pp. 197–226.

Sagratella, Simone, Marcel Schmidt, and Nathan Sudermann-Merx (2020). The noncooperative fixed charge transportation problem. *European Journal of Operational Research* 284 (1), pp. 373–382.

Sandholm, Tuomas, Andrew Gilpin, and Vincent Conitzer (2005). Mixed-integer programming methods for finding Nash equilibria. *Proceedings of the National Conference on Artificial Intelligence* 2, pp. 495–501.

Schaefer, Andrew J (2009). Inverse integer programming. *Optimization Letters* 3 (4), pp. 483–489.

Schroeder, Andreas and Thure Traber (2012). The economics of fast charging infrastructure for electric vehicles. *Energy Policy* 43, pp. 136–144.

Seim, Katja (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics* 37 (3), pp. 619–640.

Selten, Reinhard (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4 (1), pp. 25–55.

Serra, Daniel, Horst A Eiselt, Gilbert Laporte, and Charles S ReVelle (1999). Market capture models under various customer-choice rules. *Environment and Planning B: Planning and Design* 26 (5), pp. 741–750.

Serra, Daniel, Vladimir Marianov, and Charles S ReVelle (1992). The maximum-capture hierarchical location problem. *European Journal of Operational Research* 62 (3), pp. 363–371.

Shoettle, Brandon and Michael Sivak (2016). *The relative merits of battery-electric vehicles and fuel-cell vehicles*. Tech. rep. Last accessed 12.07.2022. Ann Arbor, Michigan: University of Michigan Transportation Research Institute. URL: `https://trid.trb.org/view/1480410`.

Shriver, Scott K and Bryan Bollinger (2022). Demand expansion and cannibalization effects from retail store entry: A structural analysis of multichannel demand. *Management Science*. Forthcoming: `https://doi.org/10.1287/mnsc.2022.4308`.

Simchi-Levi, David and Oded Berman (1988). A heuristic algorithm for the traveling salesman location problem on networks. *Operations Research* 36 (3), pp. 478–484.

Spence, Andrew M (1977). Entry, capacity, investment and oligopolistic pricing. *The Bell Journal of Economics* 8 (2), pp. 534–544.

Statista (2019a). Warum haben Sie Ihren letzten Einkauf gerade hier, also bei Aldi getätigt? (Why did you make your last purchase here, at Aldi?) `https://de.statista.com/statistik/daten/studie/990481/umfrage/gruende-fuer-den-einkauf-bei-aldi/`. Last accessed 17.03.2022.

— (2019b). Warum haben Sie Ihren letzten Einkauf gerade hier, also bei Lidl getätigt? (Why did you make your last purchase here, at Lidl?) `https://de.statista.com/statistik/daten/studie/990476/umfrage/gruende-fuer-den-einkauf-bei-lidl/`. Last accessed 17.03.2022.

Stein, Noah D, Asuman Ozdaglar, and Pablo A Parrilo (2008). Separable and low-rank continuous games. *International Journal of Game Theory* 37 (4), pp. 475–504.

Stengel, Bernhard von (2007). Equilibrium computation for two-player games in strategic and extensive form. *Algorithmic Game Theory*. Ed. by Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Varzirani. Cambridge University Press, pp. 53–78.

Tan, Jingzi and Wei H Lin (2014). A stochastic flow capturing location and allocation model for siting electric vehicle charging stations. *17th IEEE International Conference on Intelligent Transportation Systems*, pp. 2811–2816.

Tran, Trung H and Thu BT Nguyen (2019). Alternative-fuel station network design under impact of station failures. *Annals of Operations Research* 279, pp. 151–186.

Tsai, Jung-Fa, Ming-Hua Lin, and Yi-Chung Hu (2008). Finding multiple solutions to general integer linear programs. *European Journal of Operational Research* 184 (2), pp. 802–809.

Upchurch, Christopher and Michael Kuby (2010). Comparing the p-median and flow-refueling models for locating alternative-fuel stations. *Journal of Transport Geography* 18 (6), pp. 750–758.

Upchurch, Christopher, Michael Kuby, and Seow Lim (2009). A model for location of capacitated alternative-fuel stations. *Geographical Analysis* 41 (1), pp. 85–106.

Valiant, Leslie G (1979). The complexity of computing the permanent. *Theoretical Computer Science* 8 (2), pp. 189–201.

Viktorsson, Ludvik, Jukka T Heinonen, Jon B Skulason, and Runar Unnthorsson (2017). A step towards the hydrogen economy — A life cycle cost analysis of a hydrogen refueling station. *Energies* 10 (6). Article 763.

Wagner, Harvey M and Thomson M Whitin (1958). Dynamic version of the economic lot size model. *Management Science* 50 (1), pp. 89–96.

Wang, Lizhi (2009). Cutting plane algorithms for the inverse mixed integer linear programming problem. *Operations Research Letters* 37 (2), pp. 114–116.

Wu, Fei and Ramteen Sioshansi (2017). A stochastic flow-capturing model to optimize the location of fast-charging stations with uncertain electric vehicle flows. *Transportation Research Part D: Transport and Environment* 53, pp. 354–376.

Wu, Tai H and Jen N Lin (2003). Solving the competitive discretionary service facility location problem. *European Journal of Operational Research* 144 (2), pp. 366–378.

Zakeri, Golbon, Andrew B Philpott, and David M Ryan (2000). Inexact cuts in Benders decomposition. *SIAM Journal on Optimization* 10 (3), pp. 643–657.

Zhang, Jianzhong and Chengxian Xu (2010). Inverse optimization for linearly constrained convex separable programming problems. *European Journal of Operational Research* 200 (3), pp. 671–679.

Zhang, Yue, Oded Berman, and Vedat Verter (2012). The impact of client choice on preventive healthcare facility network design. *OR Spectrum* 34 (2), pp. 349–370.

Zhu, Ting and Vishal Singh (2009). Spatial competition with endogenous location choices: An application to discount retailing. *Quantitative Marketing and Economics* 7 (1), pp. 1–35.