Getting More from Less: Understanding Airline Pricing

John Lazarev*

Department of Economics
New York University

April 2019

Abstract

The price of an airline ticket is determined by the decisions of not one but two separate departments of the same airline: the pricing department and the revenue management department. The pricing department sets a menu of fares starting many days from the actual flight. The revenue management department treats the menu as given but decides in each moment of time which part of the menu to make available for purchase and which to keep closed. The revenue management department cannot add new fares or change the prices of the existing ones. As the result, the price of a ticket for a given flight can take only a limited number of values, not more than the number of fares in the menu. The paper offers a novel explanation for this seemingly unconventional organizational structure. It demonstrates that, in a strategic environment, firms can benefit from restricting the set of available actions as long as they can credibly commit to doing so. Creating a separate department that restricts the set of prices by choosing a menu of fares can be a way to establish and maintain the credibility of such a commitment.

*This paper is a revised chapter of my Stanford Ph.D. dissertation. I thank Peter Reiss and Andy Skrzypacz for their invaluable guidance and advice. I am grateful to Lanier Benkard, Jeremy Bulow, Alex Frankel, Ben Golub, Michael Harrison, Shota Ichihashi, Jon Levin, Trevor Martin, Shunya Noda, Michael Ostrovsky, Bob Wilson, Masayuki Yagasaki, and participants of the Stanford Structural IO lunch seminar for helpful comments and discussions. All remaining errors are my own. Correspondence: jlazarev@nyu.edu

1 Introduction

It is well known that pricing in the airline industry is complex. What is less known is that at any given moment, the price of a flight ticket is determined by the decisions of not one but two airline departments, the pricing department and the revenue management department. The pricing department sets a menu of fares starting many days from the actual flight. This menu is subsequently updated very rarely. The revenue management department treats the menu as given but decides in each moment of time which part of the menu to make available for purchase and which to keep closed. The revenue management department cannot add new fares or change the prices of the existing ones. As the result, the price of a ticket for a given flight can take only a limited number of values, not more than the number of fares in the menu. Why would airlines introduce such a complex structure, in which they effectively restrict their flexibility in changing the prices over time? At first thought it might seem they can do better by just changing the price over time without precommitting to a set of fares. One answer may be that, as long as the fare grid is fine enough, airlines do not lose much from significantly restricting the set of possible prices (Wilson, 1993). This paper offers an alternative explanation. It demonstrates that, in a strategic environment, firms can benefit from restricting the set of available actions as long as they can credibly commit to doing so. Creating a separate department that restricts the set of prices by choosing a menu of fares might be a way to establish and maintain the credibility of such a commitment.

The idea that the ability to commit to a specific action can be beneficial for an agent is a classic observation of game theory (Schelling, 1960). The contribution of this paper is to show that in a broader set of situations, agents can benefit if they are able commit to a *menu* of actions. As a simple theoretical example of why menus may be desirable, consider a Bertand game between two firms that sell identical products. The costs of production are zero. The firm that charges the lower price gets the market demand. If the prices are the same, the market demand is evenly divided between the firms. In this example, the equilibrium payoffs are zero. The ability to commit cannot increase the profit of an individual firm. If it commits to a price other than zero, the competitor can steal the market by charging a slightly lower price and getting the entire market. Thus, in this example, commitment power alone is not enough to sustain supracompetitive prices.

Consider now the following modification. The game consists of two stages: the commitment stage and the action stage. At the commitment stage, firms can simultaneously and independently decide to constrain themselves to a menu of prices. Their choices are then publicly observed. Then, at the action stage, firms independently and simultaneously choose prices from their menus. A price outside the restricted menu cannot be chosen by firms. The profits are determined according to the Bertrand game described above.

Unlike in the original game, the equilibrium payoffs in this modification can be positive. Moreover, there exists a subgame perfect equilibrium in which each firm gets a half of the monopoly profit. The following pair of symmetric strategies form this equilibrium. At the commitment stage, each firm chooses a menu that consists of two prices, the monopoly price and zero. If both firms chose these menus, then charging the monopoly price is an equilibrium of this subgame. Indeed, the only deviation at the action stage is charging zero, which reduces the profit from the half of the monopoly profit to zero. There are many deviations at the commitment stage but to be profitable they have to include a price that is higher than zero and less than the monopoly price. If a firm attempts such deviation, then the competitor will be indifferent between charging zero or the monopoly price at the action stage. If it charges zero, then the deviator will end up receiving zero. Thus, deviations at the commitment stage cannot be profitable either.

Thus, as we can see, Bertrand competitors can still achieve the monopoly outcome without signing an enforceable contract or using repeated interactions to enforce monopoly pricing. What they need to do is to exclude a set of profitable deviations from the monopoly outcome ("temptation") *and* keep "punishment" actions in case their competitors misbehave at the commitment stage.

This example is in a certain sense limited. Even when a firm undercuts by a very small amount, the profit of its competitor decreases to zero. This discontinuity of the payoffs guarantees that the punishment by zero price is a credible threat. If the competitors deviate even by a small amount, the profit of the opponent drops to zero. For continuous payoffs, these pair of strategies would not be an equilibrium anymore. This paper focuses on a class of games with continuous payoffs and characterizes pure-strategy subgame-perfect Nash equilibria of in a game that has two stages described above.

The two-stage construction of the paper has four crucial components. First, the agents are free to choose any subsets of the unrestricted action spaces. In particular, these subsets can be non-convex or include isolated actions. Non-convex subsets, as we will see later, allow players to achieve payoffs that dominate the Nash equilibrium payoffs of the original game. Second, it is important that the players are able to commit not to play actions outside of the chosen subsets. Without such commitment, the initial stage will not change the incentives of the players. Third, the chosen restricted set must become public knowledge. Finally, the players choose these subsets simultaneously. If players move sequentially, then some outcomes may not be an equilibrium to the sequential-move game.

Pricing in the airline industry has all these four components. Airline fares are discrete. The separation of the two departments creates some commitment power. Indeed, according to industry insiders, these departments do not actively interact with each other. Next, it turns out that sixteen major airlines own a company, the Airline Tariff Publishing Company (ATPCO), whose role is to collect menus of fares (but not their real-time availability) from more than 500 airlines and distribute them four times a day to all airlines, travel agents, and reservation systems. Thus, the information about the chosen subsets of fares becomes public and the decisions are made simultaneously. Therefore, all necessary features implied by the commitment concept are present in the pricing competition among U.S. airlines.

The paper has three main theoretical results. The first results shows that to support other than Nash equilibrium outcome, players have to constrain their action space. Moreover, at least one player has to choose a subset with several actions at the commitment stage. The second result provides a necessary and sufficient condition for an outcome to be supported by a commitment equilibrium. An important corollary from this result shows that in a Bertrand oligopoly firms may mutually gain from self-restraint while in Cournot they cannot. The third result demonstrates that there is a whole set of outcomes that can be supported by an equilibrium of the two-stage model. This set, among others, includes Pareto-efficient outcomes. The proof of the third result is constructive.

The history of airline pricing provides several pieces of evidence that supports the theory advanced by this paper. For example, a Harvard Business School case study known as American Airlines Value Pricing (1992) confirms the theoretical prediction of the model, namely the first result of the paper. In April 1992, American Airlines effectively tried to abandon revenue management system. They thought that fares were too complex, so the idea was to have one fare that reflects the "true value." In terms of the model, they tried to commit to a set that included only one action. Within a week, most major carriers (United, Delta, Continental, Northwest) adopted the same pricing structure. A fare war immediately followed. That is what exactly the model predicts: if players do not include any punishment actions, a Bertrand-type price war (the unique static Nash equilibrium) is the only subgame equilibrium outcome. By November 1992, American Airlines acknowledged that the plan had "clearly failed" and decided to come "back to setting and manipulating thousands of fares throughout the system."

Importantly, the goal of the paper is *not* to develop a comprehensive model of airline pricing that can be taken directly to data. Instead, the ultimate goal of this paper is to propose an explanation of one particular aspect of airline pricing, namely the presence of two, somewhat independent departments. Therefore, the approach taken here is to abstract away from many important aspects of the industry (such as intertemporal price discrimination, dynamic capacity allocation, repeated interactions, multimarket contact, multiproduct firms) and find the simple possible framework that demonstrates the main insight of the paper. This paper is one of the first papers that establish a link between the organizational structure of a firm and its market performance in an imperfectly competitive industry.

The rest of the paper is organized as follows. Section 2 reviews related literatures. Section 3 gives a brief overview of airline pricing and its history. It motivates the theoretical framework developed in the subsequent sections. Section 4 presents notations and definitions, and shows that the two-stage modification is a coarsening concept and can have as an equilibrium both Pareto better and Pareto worse outcomes compared to the Nash equilibrium outcome. In Section 5, I show that supermodular games is a natural class in which we can study the two-stage modification as it is the class in which a (pure-strategy) subgame perfect equilibrium is guaranteed to exist. Section 6 shows that there is a nontrivial set of outcomes that can be supported as an equilibrium in the two-stage game. Section 7 concludes.

2 Related Literature

This paper contributes to several literatures. First, the paper extends our understanding of the role of commitment in strategic interactions originally developed by Schelling (1960). In his work, Schelling gives examples of how a player may benefit from reducing her flexibility. In these examples, a player receives a strategic advantage by moving first and committing to a particular action. No matter what the other players do, she will not play a different action. In contrast, in this paper, the players move simultaneously. As a result, they must recognize that they need some flexibility in their actions in case other players do not restrict their actions or otherwise deviate. For example, if a player commits to a single action, then she will have no punishment should other players deviate. When a player commits to a subset of actions, then her punishment action may differ from her reward action. Moreover, the punishment may depend on the exact deviation chosen by an opponent.

At the same time, punishments have to be credible. Therefore, players do not want to include too many actions as their potential punishments. A punishment is effective only in the case when those who are supposed to punish will have an incentive to execute this punishment. The smaller is the set of available actions, the more likely the punishment is to be executed because less alternative actions are available to the player. Thus, this paper studies the trade-off between the flexibility and the credibility of available punishments. Committing to a single action is one of the extremes in this trade-off, when punishment is fully credible but completely inflexible.

Ideally, players would like to sign an enforceable contract that specifies what action each agent will play. Of course, it may be hard to specify exactly what action each agent should play. Hart and Moore (2004) recognize this fact and view a contract as a mutual commitment not to play outcomes that are ruled out by the signed contract. Bernheim and Whinston (1998) show that optimal contracts, in fact, have to be incomplete and include more than one observable outcome if some aspects of performance cannot be verified. Both papers view contracts as a means to constrain mutually the set of available actions. This paper assumes that players cannot sign such contracts. If they choose to restrict their set of available actions, they can only do it independently of each other.

By maintaining independence, this paper is more similar in the spirit to the literature that models tacit collusion. It is well known that if players interact with each other during several periods, they can achieve a certain level of cooperation even if they act independently of each other. In finite-horizon games, deviations can be deterred by a threat to switch to a worse equilibrium in later periods (Benoit and Krishna, 1985). Similarly, in infinite-horizon games, players can use a significantly lower continuation value as a punishment that supports an equilibrium in the subgame induced by a deviation. This result, known as the Folk theorem, states that if players are patient enough any individually rational outcome can be supported by a subgame perfect Nash equilibrium (see Abreu et al. (1990) and Fudenberg and Maskin (1986), among others). In contrast, this paper assumes that the game is played only once. As a result, players cannot use future outcomes to punish deviations. Instead, they strategically choose a subset of actions that must include credible punishments sufficient to deter deviations from the proposed equilibrium. These punishments must be executed immediately to affect any cooperation.

Fershtman and Judd (1987) and Fershtman et al. (1991) studied delegation games in which players can modify their payoff functions by signing contracts with agents who act on their behalf. By doing so, the players can establish commitment to play an action other than a Nash equilibrium strategy, which may result in achieving a Pareto efficient outcome. This paper takes a different approach. Instead of modifying their payoff functions, players can modify their action sets in a very specific way. They can exclude a subset of actions but cannot include anything else. Thus, the model of this paper may be viewed as a specific case of payoff modification: the players can assign a large negative value to a subset of actions but cannot modify the payoffs in any other way.

The closest papers to this paper are Bade et al. (2009) and Renou (2009). They develop a similar two-stage construction in which commitment stage is followed by a one-shot game. The constructions of their paper, however, are different in several important aspects. In the first paper, players can only commit to a *convex* subset of actions. This paper, however, allows players to choose any subset of actions. If players can choose only a convex subset of actions, then they often cannot get rid of temptations and keep actions that can be used as punishments, which are key to the results of this paper. In the homogeneous Bertrand example, supracompetitive prices cannot be sustained in equilibrium if the subsets of actions are restricted to be convex at the commitment stage.The second paper considers finite games, while this paper studies supermodular games without making any restrictions on the number of available actions.

There are a number of papers that endogenize players' commitment opportunities.[1] This paper does not address this question. The ability of firms to voluntarily restrict their action space is assumed. This assumption, however, is motivated by real world practices used by firms in an important industry: airlines.

3 Pricing in the airline industry

Back in the 1960s, buying an airline ticket was a simple process. A traveler told her travel agents when and where she wanted to fly. The travel agent then consulted a thick book that had all schedules and corresponding airfares for that season. The final step was to call the airline to verify that seats were still available and if so to book a seat on the flight. Back then airfares were regulated. A federal government agency, the Civil Aeronautics Board (CAB), used to decide which airline flew where and how much they were allowed to charge. The fares were calculated using a nonlinear function of the route's distance. They intended to cover the airlines' projected costs.

The industry changed dramatically in the 1970s when the Airline Deregulation Act was signed into law. The new law removed government control over routes, fares, or entry of new airlines. Borenstein (2004) identifies the following two "most important and unforeseen developments in the industry following deregulation." First, the development of hub-and-spoke networks. Second, airlines' increased sophistication in pricing and marketing their products. The latter development is what motivates our theoretical analysis.

The current pricing system is often characterized as being sophisticated and complex. An airline can start selling tickets on a scheduled flight as early as 330 days before departure. At any given moment, the price of a ticket is determined by the decisions of *two* airline departments, the pricing department and the revenue management department. The pricing department moves first and develops a discrete set of fares for each market. The revenue management department moves second and chooses which of the fares from this set to offer for sale at any period of time.

A fare typically contains the following elements: a fare basis code (a "name" of the fare that includes its booking class, or bucket), the origin and destination airports, permitted routing, the price, first and last ticket dates, first and last travel dates, and any restrictions

[1] See Rosenthal (1991), Van Damme and Hurkens (1996), and Caruana and Einav (2008), among others.

(e.g. minimum stay, advance purchase requirement, blackout dates, cancellation and change penalties). Surprisingly, there is little variation in these restrictions. Most fares are either unrestricted (fully refundable) or restricted (non-refundable). The non-refundable fares have the same penalties and restrictions regardless of their price, route, or how far in advance they were bought. What differentiates one restricted fare from another is mainly their price and sometimes advance purchase requirements (APR).[2] Multiple fares typically correspond to the same APR. Thus, airlines indeed offer a menu of prices for exactly the same product.

The fact that the pricing department has published a fare does not guarantee that a traveler will be able to get that fare on the specific flight. The flight needs to have available seats in the booking class that corresponds to that fare. How many seats to assign to each booking class on each flight is the primary decision of the second airline department – the revenue management department. These decisions are updated in real time. The lowest price of a ticket for a given flight is determined by the lowest fare that available seats in the corresponding booking class.

The menu of available fares is known to all airlines. Competing airlines monitor each other very closely. Three to four times a day, at predetermined times, every airline reports their set of fares to a central clearinghouse called Airline Tariff Publishing Company (ATPCO). This clearinghouse is jointly owned by all major airlines. It then produces a list of all industry fare information and sends it to every major airline and every computer reservation system.

Once the menus of available fares are transmitted, every airline knows what options their competitors' revenue management departments have. They do not know for sure which fare will be made available for sale in any given time, but they know that their competitors cannot offer fares outside of these menus.

These peculiar characteristics of airline pricing have drawn attention both from antitrust authorities and airline executives. The following two episodes from the history of airline pricing illustrate this fact. Incidentally, both episodes are consistent with the prediction of the theoretical analysis developed in the next sections of this paper.

The first episode is a famous antitrust case known in the literature as the ATPCO case of 1992 (see Borenstein (2004) for a detailed analysis). The government charged that the airlines, through ATPCO, had colluded to raise price and restrict competition. According

[2] An advance purchase requirement (e.g. a fare has to be ticketed twenty-one or more days before departure) is a screening tool designed to separate different types of customers by incentivizing more price sensitive customers to buy early. Lazarev (2018) is an empirical study that quantifies the welfare effects of intertemporal price discrimination. The focus of this paper, however, is not on advance purchase requirements or intertemporal price discrimination. The paper analyzes one-shot games. Therefore, the pricing problem can be viewed as a sequence of perhaps interdependent market segments, each bounded by the corresponding APR. I do not rely or take into account this interdependence. Instead, the theoretical results of the paper can be applied to each market segment separately.

to the complaint, ATPCO became a platform where the airlines could communicate their decisions by filing fares that will take effect only at some point of time in the future. The complaint cited several examples in which one airline had announced a fare increase only to roll it back when no other competitor matched the increased price. In other words, ATPCO was accused of being a communication device that facilitated collusion among the airlines that resulted in supra competitive prices, as the folk theorem would have predicted. The airlines denied the accusations but eventually settled agreeing not to pre-announce future price changes through ATPCO. Importantly, the settlement did not restrict any other aspect of airline pricing described above. The theoretical analysis presented in this paper shows that airlines would be able to coordinate on higher prices even without repeated interactions. This conclusion would suggest that the settlement as entered should have minimal effect on the airline prices. Miller (2010) conducted a retrospective empirical analysis and concluded that the ATPCO case "had at best a temporary effect on airline collusion." There is no evidence that the prices fundamentally changed after the settlement entered. This is exactly what the model developed in this paper would have predicted.

The second episode that provides supporting evidence in favor of the theory developed in this paper is a Harvard Business School case known as American Airlines' Value Pricing (Silk and Michael, 1993). In April 1992, the chairman, president, and CEO of American Airlines Robert L. Crandall had it enough with complexity of airline pricing.[3] He declared that airline pricing was so "complex that our customers neither understand it nor think it is fair." He dramatically proposed to reduce the total number of different fares offered by American. On any given flight, there would be only four different fares: first, coach, coach with 7-day advance purchase, and coach with 21-day advance purchase. Effectively, the menus of fares were eliminated. Proposition 1 from Section 4.3 of this paper gives some particularly good insight on what was likely to happen after this initiative was announced. Without publicly available menus, it is impossible to sustain coordination with supra competitive prices in a static game. The reaction of American Airlines' competitors was fully consistent with this prediction of the paper. Within a day or two following American's announcement, the new fare structure was adopted by most major carriers including America West, Continental, Delta, Northwest, United and USAir. Then, three days after American's initiative, TWA responded with a set of price changes that undercut Americans new regular coach fares by 10-20%. That started an industry-wide price war that lasted several months. In the wake of massive operating losses, American abandoned its Value Pricing structure in the beginning of

[3]Admittedly, this paper views these decisions of the CEO as an exogenous shock than should never have happened on the equilibrium path. Once it happened, the response of the competitors was fully consistent with their equilibrium strategies.

October 1992. In an interview with the British Broadcasting Corporation (BBC), Crandall summarized Americans view on the termination of Value Pricing: "We tried to provide some price leadership, but it didn't work, so we are back into the death by a thousand cuts."

4 Elements of a commitment equilibrium

4.1 Definitions

The game G Economic agents play a one-shot normal-form game, G. The game has three elements: a set of players $\mathcal{I} = \{1, 2, ..., n\}$, a collection of action spaces $\{\mathcal{A}_i\}_{i \in \mathcal{I}}$, and a collection of payoff functions $\{\pi_i : \mathcal{A}_1 \times \mathcal{A}_2 \times ... \times \mathcal{A}_n \longrightarrow \mathbb{R}\}_{i \in \mathcal{I}}$. Thus, $G = (\mathcal{I}, \mathcal{A}, \pi)$, where $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times ... \times \mathcal{A}_n$ and $\pi = \pi_1 \times \pi_2 \times ... \times \pi_n$. (In general, the omission of a subscript indicates the cross product over all players. Subscript $-i$ denotes the cross product over all players excluding i). The agents make decisions simultaneously and independently, which determines their payoffs. I refer to $a \in \mathcal{A}$ as an *outcome* of G and to $\pi(a) \in \mathbb{R}^n$ as a *payoff* of G.

I assume that for every player i: a) $\mathcal{A}_i$ is a compact subset of $\mathbb{R}$, and b) π_i is upper semi-continuous in a_i, a_{-i}.

An outcome $a^* \in \mathcal{A}$ is called a *Nash equilibrium (NE) outcome* if there exists a pure-strategy Nash equilibrium that supports $a^* \in \mathcal{A}$. In other words, $a^* \in \mathcal{A}$ is a *NE outcome* if and only if for all $i \in \mathcal{I}$ the following inequality holds:

$$\pi_i\left(a_i^*, a_{-i}^*\right) \geq \pi_i\left(a_i, a_{-i}^*\right) \text{ for any } a_i \in \mathcal{A}_i.$$

Denote the set of all NE outcomes by $\mathcal{E}_G$ and the set of corresponding *NE payoffs* by Π_G.

The two-stage game $C(G)$ Consider the following modification of game G. Define $C(G)$ as the following two-stage game:

Stage 0 (commitment stage). Each agent $i \in \mathcal{I}$ simultaneously and independently chooses a non-empty compact subset A_i of her action space $\mathcal{A}_i$. Once chosen, the subsets are publicly observed.

Stage 1 (action stage). The agents play game $G_C = (\mathcal{I}, A, \pi)$ where $A = A_1 \times A_2 \times ... \times A_n$. In other words, the players can choose actions that only belong to the subsets chosen at stage 0. The actions outside the chosen subsets A_i can not be played.

In this paper, I study pure-strategy subgame perfect Nash equilibria of $C(G)$.

A *strategy* for player i in the game $C(G)$ is a set A_i and a function σ_i which selects, for any subsets of actions chosen by players other than i, an element of A_i, i.e. $\sigma_i : \mathbb{A}_{-i} \longrightarrow A_i$. A pure-strategy subgame perfect Nash equilibria of $C(G)$ is called an independent simultaneous self-restraint (commitment) equilibrium of the game G. Formally, *a commitment equilibrium* is an n-tuple of strategies, (A^*, σ^*), such that

(i) for all i and any set of actions $\widetilde{A}_i \in \mathbb{A}_i$,

$$\pi_i\left(\sigma^*\left(A^*\right)\right) \geq \pi_i\left(\widetilde{\sigma}_1^*, ..., \widetilde{\sigma}_{i-1}^*, \widetilde{a}_i, \widetilde{\sigma}_{i+1}^*, ..., \widetilde{\sigma}_n^*\right) \text{ for any } \widetilde{a}_i \in \widetilde{A}_i,$$

where $\sigma^*\left(A^*\right) = \left(\sigma_1^*\left(A_{-1}^*\right), \ldots, \sigma_n^*\left(A_{-n}^*\right)\right)$, $\widetilde{\sigma}_j^* = \sigma_j^*\left(\widetilde{A}_{-j}^*\right)$ and $\widetilde{A}_{-j}^* = A_1^* \times ...A_{j-1}^* \times A_{j+1}^* \times ...A_{i-1}^* \times \widetilde{A}_i \times A_{i+1}^* ... \times A_n^*$;

(ii) for all i and any $\widetilde{A}_i$, there exists $\widetilde{a}_i \in \widetilde{A}_i$, such that $\left(\widetilde{a}_i, \widetilde{\sigma}_{-i}^*\right)$ is a pure strategy Nash equilibrium of game $\widetilde{G} = \left(\mathcal{I}, A_1^* \times ...A_{i-1}^* \times \widetilde{A}_i \times A_{i+1}^* ... \times A_n^*, \pi\right)$.

An outcome $a^* \in \mathcal{A}$ is called a *commitment equilibrium outcome* if there exists a commitment equilibrium that supports $a^* \in \mathcal{A}$, i.e. $a^* = \left(\sigma_1^*\left(A_{-1}^*\right), ..., \sigma_n^*\left(A_{-n}^*\right)\right)$. Denote the set of all commitment equilibrium outcomes by $\mathcal{E}_C$ and the set of corresponding *commitment equilibrium payoffs* by Π_C.

4.2 Reward, temptation, and punishment

To better describe the structure of a commitment equilibrium, I introduce the concepts of reward, temptation and punishment actions. Take any commitment equilibrium, let $a_i^* = \sigma_i^*\left(A_{-i}^*\right)$ denote the *reward* action of player i.

For an outcome $a^* \in A_i$, an action $a_i^T \in A_i$ is called a *temptation in* A_i of player i if $\pi_i\left(a_i^T, a_{-i}^*\right) > \pi_i\left(a_i^*, a_{-i}^*\right)$. Define the set of temptations $T_i\left(a^*|A_i\right)$ as:

$$T_i\left(a^*|A_i\right) = \left\{a_i^T \in A_i : \pi_i\left(a_i^T, a_{-i}^*\right) > \pi_i\left(a_i^*, a_{-i}^*\right)\right\}.$$

It follows from the definition of commitment equilibrium that if a^* is an outcome supported by a commitment equilibrium (A^*, σ^*), then none of the players has a temptation in $A_i^* \subseteq \mathcal{A}_i$. In other words, $T_i\left(a^*|A_i^*\right) = \varnothing$ for all i. Obviously, a^* is a Nash equilibrium outcome if and only if none of the players has a temptation in $\mathcal{A}_i$.

A commitment equilibrium can include not only actions that are played along the equilibrium path but also punishment actions that deter profitable deviations from the intended outcome. Let $A_i^P = A_i^* \backslash \left\{a_i^*\right\}$ and $a_i^P \in A_i^P$ denote a *punishment set* and a *punishment action*, respectively. Depending on the payoff function and the outcome that the equilibrium intends to support, players may need to use different punishments against different deviations. However, even one punishment may give the players a powerful device that allows them to coordinate on better outcomes. Of course, A_i^P may be empty for one or all players. In the latter case, a commitment equilibrium outcome will coincide with a NE outcome, as Proposition 1 will demonstrate.

Therefore, the intuition behind commitment equilibria is the following. On the one hand, players want to exclude all temptations from their action space when they choose to restrain their sets of available actions. On the other hand, since they cannot promise to play the reward actions at the action stage, the players need to include some actions that can be used as punishments should other players deviate. To determine under what conditions the chosen punishments can deter other players from deviating, we need to define and analyze the subgames induced by players' choices at the commitment stage.

4.3 Relation to Nash equilibrium

I now show that any NE outcome can be supported by a commitment equilibrium. Take any Nash equilibrium. Suppose that at the commitment stage all players choose their NE action, excluding all other actions from their sets. At the action stage, there are no profitable deviations since there is only one action available to play. For this to be a commitment equilibrium, we need to show that there are no deviations available at the commitment stage. But this result follows almost immediately. Indeed, if all but one player chose their NE action, then the remaining player cannot choose any other subset and find it profitable to deviate by doing so. To be complete, I must show that the deviator's payoff attains its maximum on any subset. This requires assuming that the payoff functions are upper semi-continuous and the action subsets are compact, which was a part of the construction's definition.

Commitment equilibrium is a coarsening concept. Any Nash equilibrium outcome can be supported by a commitment equilibrium in which players restrict their subsets of actions at the commitment stage to a singleton. If one player fails to do so and leaves her action set unconstrained, the other player might find it beneficial to constrain himself to playing his Stackelberg action, i.e. the action that maximizes his payoff subject to his opponent's best response.

The main question of this paper, however, is whether commitment equilibria can achieve better (or worse) payoffs compared to NE outcomes. It turns out that the number of actions that the players keep at the commitment stage is a key factor that determines what outcomes can be supported in a commitment equilibrium. It is not sufficient to exclude deviations that will be profitable. Players have to have the ability to carry out punishments in order to motivate other players not to deviate at the commitment stage. The next proposition formalizes this idea.

Proposition 1 ***(To get a carrot, one needs to publicly carry a stick)*** *Suppose at stage* 0 *players can only choose subsets* A_i *that include only one element, i.e.* $\mathbb{A}_i = \{A_i : |A_i| = 1\}$. *Then commitment at stage* 0 *does not produce new equilibrium outcomes, i.e.* $\mathcal{E}_G = \mathcal{E}_C$.

Proof. The proof shows that if a^0 is not a NE outcome, then it cannot be a commitment equilibrium outcome in the case when $|A_i| = 1$. If a^0 is not a NE, then there exists a player j and an action a'_j such that $\pi_j\left(a^0_j, a^0_{-j}\right) < \pi_j\left(a'_j, a^0_{-j}\right)$. Suppose a^0 is in fact a commitment equilibrium. Since for all i, $|A_i| = 1$, it holds that $\sigma^0_i\left(\left\{a^0_j\right\}, \left\{a^0_{-j}\right\}\right) = \sigma^0_i\left(\left\{a'_j\right\}, \left\{a^0_{-j}\right\}\right) = a^0_i$. But then

$$\pi_j\left(a^0\right) = \pi_j\left(a^0_j, a^0_{-j}\right) < \pi_j\left(a'_j, a^0_{-j}\right) = \pi_j\left(\sigma^0_1, ..., \sigma^0_{j-1}, a'_j, \sigma^0_{j+1}, ..., \sigma^0_n\right),$$

which violates the first condition of commitment equilibrium. Thus, $\overline{\mathcal{E}}_G \subseteq \overline{\mathcal{E}}_C$. ■

In other words, to support an outcome outside $\mathcal{E}_C$, players have to choose more than one action at the commitment stage. If players want to achieve payoffs outside the NE set, they have to constrain themselves, but not too much.

4.4 Examples

The following example shows that it is in fact possible to support an outcome that Pareto dominates the NE one.

Example 1 (Commitment equilibrium that Pareto dominates NE) *Consider the following two-by-two game:*

	C1	*C2*
R1	$1, 1$	$-1, -1$
R2	$2, -1$	$0, 0$

Row player has a strictly dominant strategy R2. Column player will respond by playing C2. Thus, (R2,C2) is the only NE outcome. However, outcome (R1,C1) can be supported by a commitment equilibrium. Indeed, let $A_1 =$ *(R1),* $A_2 =$ *(C1,C2). At stage* 1*, Row player plays R1 and Column player plays C1 if Row player plays R1. If Row player deviates at stage* 0*, Column player punishes with C2. It is easy to verify that this set of strategies is a commitment equilibrium.*

Two things are important in Example 1. First, Row player commits not to use his dominant action R2. Second, Column player has an ability to punish Row player by playing

C2 if she decides to deviate. This punishment must be included at stage 0 and this fact must be publicly known.

In Example 1, the commitment equilibrium outcome is better for both players than the unique Nash equilibrium of the original game. The next example shows that the opposite can be also true: an outcome that is strictly worse for both players than the unique Nash equilibrium can be supported by a commitment equilibrium.

Example 2 (Commitment equilibrium that is Pareto dominated by NE) *Consider the following three-by-three game:*

	C1	*C2*	*C3*
R1	$2, 2$	$2, -1$	$0, 0$
R2	$-1, 2$	$1, 1$	$-1, -1$
R3	$0, 0$	$-1, -1$	$-1, -1$

Row player has a strictly dominant strategy R1. Column player has a strictly dominant strategy C1. Thus, (R1,C1) is the unique NE outcome. However, outcome (R2,C2) can be supported by a commitment equilibrium. Let $A_1 = (R2,R3)$, $A_2 = (C2,C3)$. *At stage* 1, *Row player plays R2 and Column player plays C2 if the opponents chose the equilibrium subsets. They play R3 or C3 otherwise. This set of strategies is a commitment equilibrium that leads to an outcome that is strictly worse than the unique Nash equilibrium outcome for both players.*

The next example shows that in some cases a player has to have several punishments to achieve a better equilibrium outcome.

Example 3 (To get the carrot, one may need several sticks) *Consider the following three-by-three game:*

	C1	*C2*	*C3*
R1	$0, 0$	$0, -1$	$-1, 0$
R2	$-1, 3$	$1, 1$	$0, 2$
R3	$-1, 3$	$0, -2$	$2, -1$

Column player has a dominant strategy C1. Row player will respond by playing R1. Thus, (R1,C1) is the unique NE outcome. However, outcome (R2,C2) can be supported by a commitment equilibrium. Let $A_1 = (R1,R2,R3)$, $A_2 = (C2)$. *At stage* 1, *Row player plays R2 if Column player chose* A_2, *R3 if Column player chose C3 or* $\{C2,C3\}$ *and R1 otherwise. Column player plays C2. This set of strategies is a commitment equilibrium. Note that (R2,C2)* *can be supported by an equilibrium only when Row player is able to include all three actions at stage* 0*.*

Thus, to support an outcome outside the NE set, players must include at least two actions in their subsets at the commitment stage. In some cases, players have to include multiple punishments. One punishment may be effective and credible against one deviation, another may be effective and credible against other deviations.

5 Subgame supermodular games

5.1 Definitions

Proposition 1 showed that to support an outcome outside of the set of NE outcomes, at least one of the players has to announce more than one action at the commitment stage. Since some other player can deviate to a subset that includes more than one action, a subgame induced by such deviation, in principle, may not possess a pure-strategy Nash equilibrium, which will violate the second requirement of a commitment equilibrium. Thus, to guarantee the existence of a commitment equilibrium that can support other than NE outcomes, we need to make sure that a pure-strategy Nash equilibrium exists in any subgame following the commitment stage.

There are two approaches in the literature that establish the existence of pure-strategy equilibria. The first approach derives from the theorem of Nash (1950). The conditions of the theorem require the action sets be nonempty, convex and compact, and the payoff functions to be continuous in actions of all players and quasiconcave in its own actions. Reny (1999) relaxed the assumption of continuity by introducing an additional condition on the payoff functions known as better-reply security. The second approach was introduced by Topkis (1979) and further developed by Vives (1990) and Milgrom and Roberts (1990),[4] among others. They proved that any supermodular game has at least one pure-strategy Nash equilibrium.

The first approach requires the action space be convex, while the second approach places no such restrictions. Even if the unconstrained action space was convex, players could choose a non-convex subset at the commitment stage. As the result, there may exist a subgame that has no pure-strategy Nash equilibria. Therefore, to guarantee the existence of an equilibrium in any subgame, we will follow the second approach and focus our attention on supermodular games.

[4] Migrom and Roberts (1990) lists a number of economic models that are based on supermodular games.

Game G is called *supermodular* if for every player i, π_i has increasing differences in (a_i, a_{-i}).[5] Game $C(G)$ is called *subgame supermodular* if any subgame induced by a unilateral deviation at the commitment stage is supermodular. It is easy to show that $C(G)$ is subgame supermodular if and only if G is supermodular. First, suppose that $C(G)$ is subgame supermodular. Consider a subgame in which players did not restrict their actions in the first stage. By definition, this subgame is supermodular, therefore G is supermodular. Now, suppose that G is supermodular, consider any subgame of $C(G)$. For any non-empty compact subsets A_i the payoff functions π_i still have increasing differences (See Topkis, 1979). Therefore, $C(G)$ is subgame supermodular. The existence of a commitment equilibrium follows from the fact that game G has a pure-strategy Nash equilibrium since it is a supermodular game.

5.2 Single-action deviation principle

To prove that a collection of restricted action sets together with a profile of proposed actions form a subgame perfect equilibrium, we need to show that no profitable deviation exists. The absence of profitable deviations at the action stage is easy to establish. The only condition we need to verify is to check if the players excluded all temptations from their action sets. The commitment stage is more involved since a player can possibly deviate to any subset of the unrestricted action set. Luckily, the following result demonstrates that it is sufficient to check if there exists a profitable deviation to subsets that includes only one action.

Lemma 1 *An outcome a^* from a collection of restricted action sets A^* can be supported by a commitment equilibrium if and only if no player can profitably deviate to a singleton. Formally, for each player i and each her action a_i, there exists an equilibrium outcome $\tilde{a}$ in a subgame induced by sets $A_i = \{a_i\}$ and A^*_{-i} such that $\pi_i(\tilde{a}) \leq \pi_i(a^*)$.*

Proof. If a player can profitably deviate to a singleton at the commitment stage, then the original construction is not a commitment equilibrium. Suppose now that a player can profitably deviate to a set that includes more than one action. Then there exists an equilibrium in the subgame that is induced by the deviation that gives the player a higher payoff. Suppose instead of deviating to the set the player chooses only the action that is played in that equilibrium. It is easy to see that this outcome will still be equilibrium. This player does not have any deviations and the other players can still play the same equilibrium actions: elimination of irrelevant actions cannot generate profitable deviations

[5] A function $f : S \times T \to R$ has *increasing differences* in its two arguments (s, t) if $f(s, t) - f(s, t')$ is increasing in s for all $t \geq t'$.

in this subgame. Thus, if there exists a profitable deviation to a set, then a profitable deviation to a singleton has to exist. ■

Thus, to find commitment equilibrium outcomes, it is enough to consider deviations to singletons. To do that, it is convenient to work with best response functions defined for subgames induced by deviations at the commitment stage.

5.3 Subgame best response

Suppose at the commitment stage player i chose a subset of actions A_i. If player i follows a commitment equilibrium strategy, then she must choose an action from her subset that maximizes her payoff in the subgame. Formally, if the other players play $a_{-i} \in A_{-i}$, then her equilibrium strategy has to choose an action $\sigma_i(A_{-i}) \in BR_i(a_{-i}|A_i)$, where:

$$BR_i(a_{-i}|A_i) = Arg \max_{a_i \in A_i} \pi_i(a_i, a_{-i}).$$

I will call $BR_i(a_{-i}|A_i)$ player i's *subgame best response*. Note that $BR_i(a_{-i}|\mathcal{A}_i)$ is a standard best-response correspondence of player i in game G.

Several properties of subgame best responses follow from continuity of the payoff functions. By the Maximum Theorem of Berge (1979), the subgame best response $BR_i(a_{-i}|A_i)$ is non-empty, compact-valued and upper semi-continuous for any nonempty set A_i. If game G is supermodular, then in addition to these properties, the subgame best response is non-decreasing (Topkis, 1979).

Based on the continuity of the subgame best-response, we can analyze the relation between subgame best responses for unrestricted and restricted action sets. Denote by $BR_i^U(a_{-i})$ the subgame best response for an unrestricted set. Consider a restricted subgame best response for a set A_i.

First, suppose that A_i includes an interval $[\underline{a}_i, \overline{a}_i] \in A_i$ that belongs to the image of the unrestricted subgame best response $BR_i^U(a_{-i})$. Then for any interior point of this interval $a_i \in (\underline{a}_i, \bar{a}_i)$, the inverse images of the restricted and unrestricted subgame best responses coincide: $BR_i^{-1}(a_i\,|\mathcal{A}_i) = BR_i^{-1}(a_i\,|A_i)$. Figure 1 illustrates this property for the case of a differentiated Bertrand duopoly with linear demand functions.[6] The black line shows the unrestricted subgame best response for Firm 1, i.e. it shows what price p_1 Firm 1 should charge if Firm 2 charges price p_2. If the demand function is linear, then the unrestricted subgame best response function is linear as well. Suppose Firm 1 keeps an interval of prices $\left[p_1^L, p_1^H\right]$ in its menu. Then for any price p_1 that in the interior of this interval, the restricted and unrestricted subgame best responses coincide.

[6] See Appendix for the description of the model.

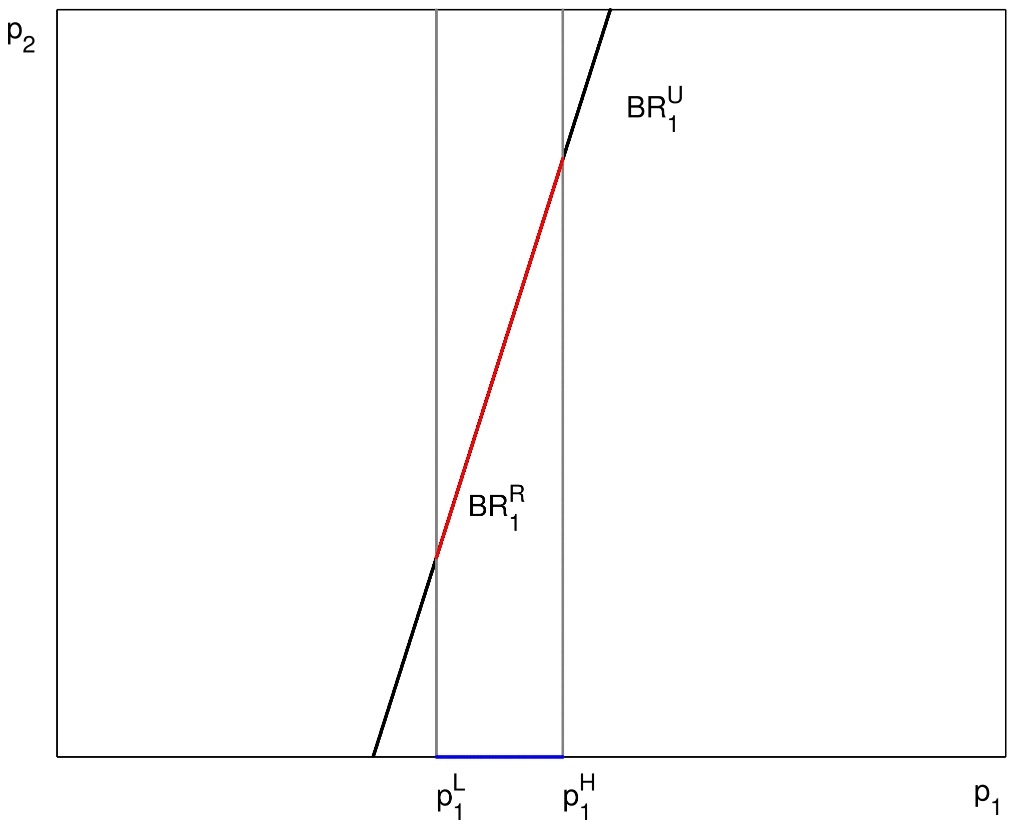

Figure 1: Firm 1's restricted and unrestricted subgame best response I

Second, suppose that two points from the image of the unconstrained subgame best response $BR_i(\mathcal{A}_{-i}\,|A_i)$ belong to the restricted set A_i but the interval between them $(\underline{a}_i, \overline{a}_i)$ does not. Then there exists a vector of the competitors' actions a_{-i} such that both points of the interval belong to the restricted subgame best response for this vector: $\{\underline{a}_i; \overline{a}_i\} \in BR_i(a_{-i}\,|A_i)$ and the agent is indifferent between playing either of them $\pi_i(\underline{a}_i, a_{-i}) = \pi_i(\overline{a}_i, a_{-i})$. Figure 2 illustrates this result for the differentiated Bertrand duopoly with linear demand functions. Suppose that Firm 1 kept two prices p_1^L and p_1^H but excluded the interval between them at the commitment stage. Then, there exists a price p_2^i such that if Firm 2 announces this price, Firm 1 will be indifferent between charging prices p_1^L and p_2^H. The isoprofit crossing points $\left(p_1^L, p_2^i\right)$ and $\left(p_1^H, p_2^i\right)$ illustrates this fact.

5.4 Subgame equilibrium response

For games with two players it is enough to analyze restricted subgame best responses. Indeed, it is enough to check that the equilibrium action of the player maximizes her profit subject to

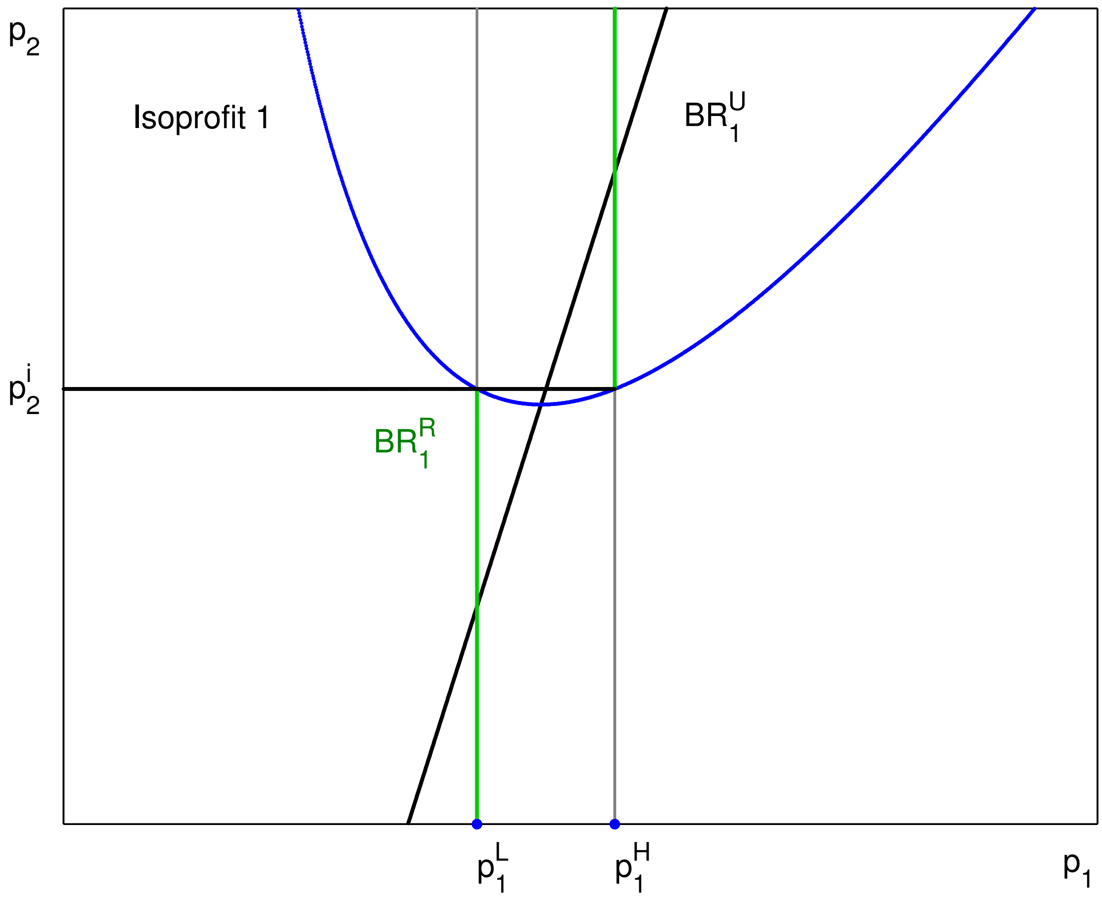

Figure 2: Firm 1's restricted and unrestricted subgame best response II

the fact that the opponent plays his restricted subgame best response. Formally, an outcome $a^* = (a_1^*, a_2^*) \in A_1 \times A_2$ can be supported by a commitment equilibrium if and only if for each player i and her action a_i, it holds that $\pi_i(a_i, \tilde{a}_{-i}) \leq \pi_i(a^*)$ where $\tilde{a}_{-i} \in BR_{-i}\left(a_i | A_{-i}^*\right)$.

To analyze games with three or more players, looking at best responses is not enough. We need to find a new equilibrium in each subgame induced by a possible deviation. Since more than one player will react to the deviation, the subgame best responses have to take into account not only the action of the deviator but also the reaction of the other players to this deviation. To take this into account, we define a subgame equilibrium response as a function that for each possible deviation of player i defines an equilibrium response of her opponents. Formally, function $ER_{-i}(a_i|A_{-i}) = a_{-i}$, where (a_i, a_{-i}) is an equilibrium of the subgame for restricted sets $A_i = \{a_i\}$ and A_{-i}. If there are several equilibria in this subgame, ER_{-i} chooses the one for which $\pi_i(a_i, a_{-i})$ is smallest. Using the subgame equilibrium response function, we can formulate a necessary and sufficient condition for a commitment equilibrium.

Proposition 2 *A collection of restricted sets A_i^*, $i = 1, ...n$ can support an outcome $a^* \in A^*$ by a commitment equilibrium if and only if for each agent i and action a_i, $\pi_i\left(a_{\,i}, ER_{-i}\left(a_i|A_{-i}^*\right)\right) \leq \pi_i\left(a_i^*, a_{-i}^*\right)$.*

Proof. If the inequality does not hold for some a_i, then choosing a_i at the commitment stage is a profitable deviation for player i. So, we need only to show that the reverse is true. By Lemma 1, it is enough to show that no profitable deviation to a singleton exists. Suppose it does and a_i is such a deviation. Then in the subgame induced by this deviation there is an equilibrium in which the opponents of the deviator play $a_{-i} = ER_{-i}\left(a_i|A_{-i}^*\right)$. But since $\pi_i\left(a_i, ER_{-i}\left(a_i|A_i^*\right)\right)$ is less than the equilibrium payoff for any a_i, this deviation cannot be profitable. ■

Proposition 2 can be restated using the notion of a superior set. For each player i define a *superior set* as $S_i(x) = \{a \in \mathcal{A} : \pi_i(a) > \pi_i(x)\}$. Then an outcome $a^* \in \mathcal{A}$ can be supported by a commitment equilibrium if and only if for all i the intersection of $S_i(a^*)$ and the graph of $ER_{-i}(a_i\ |A_{-i})$ is empty. The fact we just stated illustrates the role the commitment stage plays. For a given outcome $a \in A$, the superior set $S_i(a^*)$ is fixed but the best response equilibrium depends on the chosen restricted sets. To support a commitment equilibrium, the players have to chose a collection of restricted action sets such that the graph of the resulting equilibrium response does not overlap with the superior set.

5.5 Credibility of punishment

For a subgame supermodular game, the subgame best response is a nondecreasing correspondence. This fact implies that $ER_{-i}(a_i|A_{-i})$ is a nondecreasing function for any A_{-i}. Consider the incentives of player i to play an outcome $a \in \mathcal{A}$ as a commitment equilibrium. Suppose he decides to deviate by playing $a_i' > a_i$. This deviation will not be profitable if $\pi_i\left(a_i', ER_{-i}\left(a_i', A_{-i}\right)\right) \leq \pi_i(a)$. The same equality should hold if the deviator decides to increase its action. Thus, the payoff function cannot increase or decrease in all its arguments around the commitment equilibrium point. Since the equilibrium response function is nondecreasing, only those outcomes for which the temptation set for each player lies below the reward action, i.e. $T_i(a_i|A_i) \leq a_i$ for all i, can be supported by a commitment equilibrium.

For a Bertrand oligopoly with differentiated products this property holds if the firms want to support higher than NE prices. Indeed, when the prices exceed the NE, the profit of a potential deviator will increase if he reduces its price but decrease if his competitors decide to do so.

Cournot duopolies,[7] however, will violate this property. Suppose the firms want to support a higher than monopoly price. To do so, they need to produce less. Thus, the deviator's profit will increase if he increases his output. To punish him, the firms have to increase their output further, which will drive the price down. However, their best response to an increase in total output is to decrease their own. Thus, in Cournot oligopolies, punishments that could have supported a high price are not credible.

6 Stackelberg set

This subsection shows that there is, in fact, a set of payoff profiles that contains Pareto efficient payoffs that can be supported by a commitment equilibrium. To define this set, we need to use the concept of Stackelberg outcomes. An outcome a^{L_i} is called *a Stackelberg outcome for player i*, if $\pi_i\left(a_i^{L_i}, ER_{-i}\left(a_i^{L_i}|\mathcal{A}_{-i}\right)\right) \geq \pi_i\left(a_i, a_{-i}\right)$ for any $a_i \in \mathcal{A}_i$. This outcome corresponds to a subgame perfect equilibrium in the following modification of game G: player i moves first; the rest of the players move simultaneously and independently of each other after observing the action of player i. Similarly to a Nash equilibrium outcome, it is easy to establish that a Stackelberg outcome for any player i can be supported by a commitment equilibrium. This equilibrium would prescribe that player i should choose their equilibrium action only at the commitment stage, while the rest of the players keep their action space unrestricted. If all players do that, then none will have an incentive to deviate either at the commitment stage or at the action stage.

Define the payoff of player i that corresponds to a Stackelberg outcome by $\pi_i^{L_i}$. The Stackelberg set L is then defined as:

$$L = \left\{a \in \mathcal{A}\colon\ \pi_i\left(a\right) \geq \pi_i^L \text{ and } T_i\left(a|\mathcal{A}_i\right) \leq a_i \text{ for all } i\ \right\}.$$

The previous literature has shown that the Stackelberg set L for a class of games is not empty (Amir and Stepanova (2006) and Dowrick (1986)). Figure 3 shows the Stackelberg set for the Bertrand duopoly with differentiated products and linear demands.

The next proposition shows that any outcome from the Stackelberg set can be supported by a commitment equilibrium if the players have an action that can serve as a credible punishment. The proof of the proposition is constructive.

[7] Generically, a Cournot oligopoly with three or more firms is not a supermodular game. However, it may become one under a set of reasonable assumptions placed on the market demand and cost functions (see, Amir, 1996). The argument will still be true for games that satisfy these assumptions.

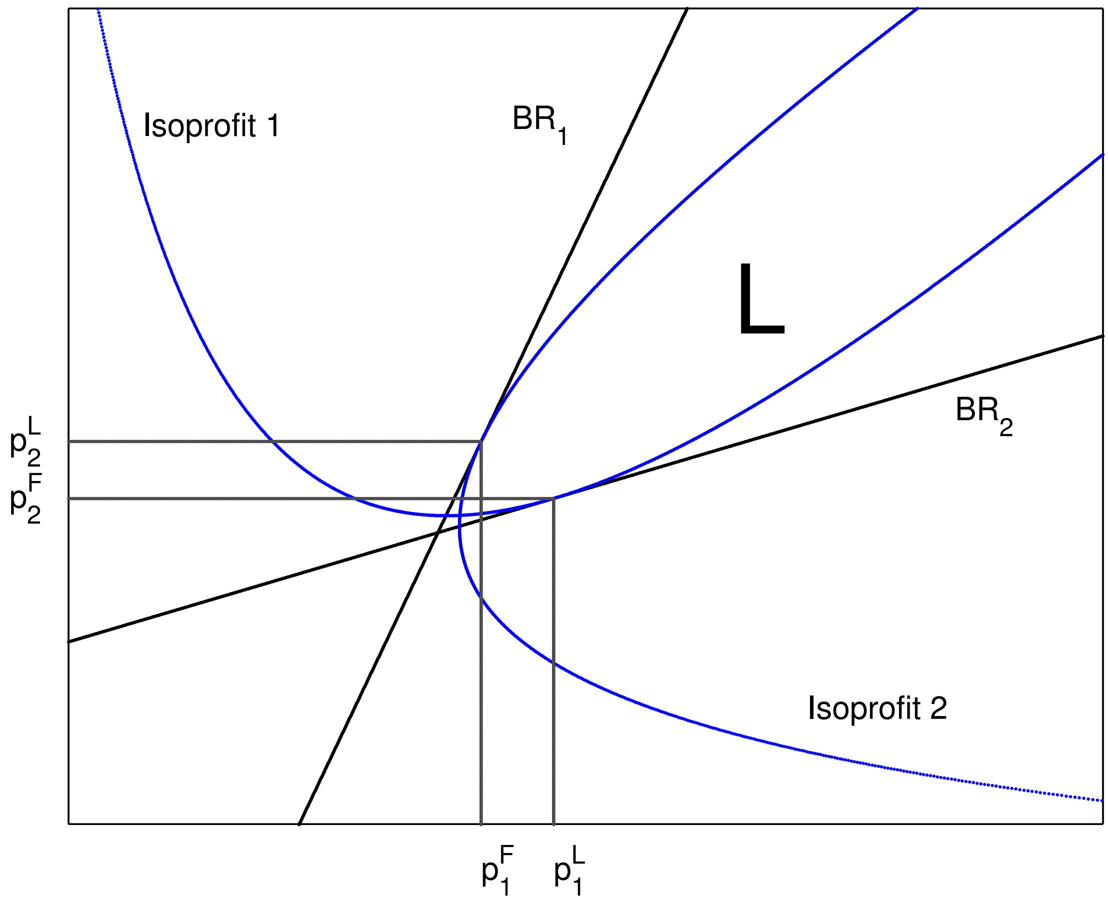

Figure 3: Stackelberg set

Proposition 3 *Suppose for each player i, the payoff function π_i is increasing in all a_{-i}. Then an outcome $a^* \in L$ can be supported by a commitment equilibrium if for any i there exists an action a_i^P satisfying $\pi_i\left(a_i^P, a_{-i}^*\right) = \pi_i\left(a_i^*, a_{-i}^*\right)$ such that $a_i^P \leq BR_i\left(a_{-i}^* \mid \mathcal{A}_i\right)$.*

Proof. Consider the following set of strategies. At the commitment stage each player includes in A_i^* the following elements: a_i^*, a_i^P and all actions below a_i^P. To prove that a^* can be supported by a commitment equilibrium with these restricted sets, it is enough to show that for each player i the graph of the equilibrium best response function and the superior set do not overlap.

First, consider the restricted subgame best responses. Based on the properties of the restricted subgame best response established in Section 4.3, we can conclude that for $a_{-i} < a_{-i}^P$, the best response functions coincide. Then the restricted best response is equal to a_i^P until a_{-i}^*. At a_{-i}^* the player is indifferent between playing a_i^P and a_i^*. Finally, for $a_{-i} > a_{-i}^*$, player i will be playing a_i^*. Thus, the equilibrium response function will be the following: for $a_i \geq a_i^*$, the other players will play a_{-i}^*, for $a_i < a_i^*$, they will play a_{-i}^P until $a_i \in BR_i\left(a_{-i}^P|\mathcal{A}_i\right)$ and $ER_{-i}\left(a_i|\mathcal{A}_{-i}\right)$ otherwise.

Since $a^* \in L$, the superior set of player i belongs to the superior set constructed for the Stackelberg outcome for player i, which in turn does not overlap with the unrestricted equilibrium response function. For $a_i < a_i^*$ the graph of the restricted equilibrium responses is bounded from above by the unrestricted equilibrium response since $a_i^P \leq BR_i\left(a_{-i}^*|\mathcal{A}_i\right)$. Therefore, we need to check that there are no overlaps for $a_i > a_i^*$. However, since the temptation set lies below a_i^*, such overlaps cannot exist. ■

The proposition shows that the punishment action has to satisfy the following property. Conditional on the fact that all the opponents play the reward action, the player has to be indifferent between playing his reward action and playing the punishment action. If a deviator follows his temptation and includes an action from the temptation set, the other players will execute the punishment, which precludes the deviator from receiving any benefits from not choosing the reward.

7 Concluding comments

Motivated by pricing practices in the airline industry, this paper shows that in a competitive environment, agents may benefit from committing not to play certain actions. For the commitment equilibrium construction to work, all players should have commitment power and be able to play several actions after the commitment stage. The main intuition of the paper is the following: to get the reward, the players need to exclude all temptations but keep punishments to motivate their opponents. Although the model is simple, its intuition can explain the otherwise puzzling menu structure of airline fares.

The airline industry plays an important role in the U.S. and world economy. It generates more than 10 million American jobs and 5 percent of the U.S. gross domestic product and nearly \$1.5 trillion in annual economic activity. Although the paper was motivated by airlines, the results developed in this paper can shed some light on pricing patterns in other industries too. For example, DellaVigna and Gentzkow (2017) documents a nearly uniform pricing in U.S. retail chains. Typically, a chain's headquarter or regional offices give their stores only a coarse grid of prices (e.g. regular price, sale price, and clearance price) and let the stores decide which one of them to choose. As a result, in the data, we see nearly uniform prices within that do not respond to local demand factors. This paper provides one of plausible explanations to this phenomenon. Retail chains can get more from giving less price flexibility to their stores.

It may seem that the theory developed in the paper is at odds with the fact that airlines use coarse pricing in monopoly markets as well. There are two possible arguments that could counter this criticism. First, separating pricing from revenue management may alleviate the Coase conjecture price dynamics associated with durable goods. Second, empirically, we do see that the lowest fares from the menus are offered significantly more frequently in monopoly markets than in duopoly markets, which is again consistent with this paper's analysis. Of course, there could be alternative explanations to this empirical observation. An in-depth empirical investigation is required but it is outside the scope of this paper.

There are several interesting extensions that could be further investigated. First, the model assumed that the action stage was played only once. It is interesting to see what the players can achieve when they can play the game a finite number of times. The commitment stage in this case may provide punishments that players execute not only after deviations at the commitment stage but also following deviations at the preceding action stages. Second, the trade-off becomes more complicated when there is some uncertainty that is resolved between the commitment and action stages. In this case, at the commitment stage the players do not know the exact punishment that they will need to include and may end up including a non-credible punishment or a temptation. Therefore, there is always a probability that the award action will not be an equilibrium at the action stage along the equilibrium path. The third extension is more technical. Renou (2009) shows an example of a game in which a mixed strategy Nash equilibrium does not survive in the two-stage commitment game. Thus, the question is which mixed strategy equilibria can and which cannot be supported by a commitment equilibrium.

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, September 1990, *58* (5), 1041–1063.

Amir, Rabah, "Cournot Oligopoly and the Theory of Supermodular Games," *Games and Economic Behavior*, August 1996, *15* (2), 132–148.

___ **and Anna Stepanova**, "Second-mover advantage and price leadership in Bertrand duopoly," *Games and Economic Behavior*, April 2006, *55* (1), 1–20.

Bade, Sophie, Guillaume Haeringer, and Ludovic Renou, "Bilateral commitment," *Journal of Economic Theory*, July 2009, *144* (4), 1817–1831.

Benoit, Jean-Pierre and Vijay Krishna, "Finitely Repeated Games," *Econometrica*, July 1985, *53* (4), 905–922.

Bernheim, B. Douglas and Michael D. Whinston, "Incomplete Contracts and Strategic Ambiguity," *The American Economic Review*, September 1998, *88* (4), 902–932.

Borenstein, Severin, "Rapid price communication and coordination: The airline tariff publishing case (1994)," *The Antitrust Revolution: Economics, Competition, and Policy*, 2004, *4*.

Caruana, Guillermo and Liran Einav, "A Theory of Endogenous Commitment," *Review of Economic Studies*, 2008, *75* (1), 99–116.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform pricing in us retail chains," 2017.

Dowrick, Steve, "von Stackelberg and Cournot Duopoly: Choosing Roles," *The RAND Journal of Economics*, July 1986, *17* (2), 251–260.

Fershtman, Chaim and Kenneth L. Judd, "Equilibrium Incentives in Oligopoly," *The American Economic Review*, December 1987, *77* (5), 927–940.

___ **,** ___ **, and Ehud Kalai**, "Observable Contracts: Strategic Delegation and Cooperation," *International Economic Review*, 1991, *32* (3), 551–559.

Fudenberg, Drew and Eric Maskin, "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, May 1986, *54* (3), 533–554.

Hart, Oliver and John Moore, “Agreeing Now to Agree Later: Contracts that Rule Out but do not Rule In,” *National Bureau of Economic Research Working Paper Series*, March 2004, *No. 10397.*

Lazarev, John, “The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in US monopoly markets,” 2018.

Milgrom, P. and J. Roberts, “Rationalizability, learning, and equilibrium in games with strategic complementarities,” *Econometrica*, 1990, *58* (6), 1255–1277.

Miller, Amalia R, “Did the airline tariff publishing case reduce collusion?,” *The Journal of Law and Economics*, 2010, *53* (3), 569–586.

Nash, J. F, “Equilibrium points in n-person games,” *Proceedings of the National Academy of Sciences of the United States of America*, 1950, pp. 48–49.

Renou, Ludovic, “Commitment games,” *Games and Economic Behavior*, May 2009, *66* (1), 488–505.

Reny, P. J, “On the existence of pure and mixed strategy Nash equilibria in discontinuous games,” *Econometrica*, 1999, *67* (5), 1029–1056.

Rosenthal, Robert W., “A note on robustness of equilibria with respect to commitment opportunities,” *Games and Economic Behavior*, May 1991, *3* (2), 237–243.

Schelling, Thomas C., *The strategy of conflict*, Harvard University Press, 1960.

Silk, A.J. and S.C. Michael, “American Airlines Value Pricing (A),” *Harvard Business School Case Study 9-594-001. Cambridge, MA: Harvard Business School Press*, 1993.

Topkis, D. M, “Equilibrium points in nonzero-sum n-person submodular games,” *SIAM Journal on Control and Optimization*, 1979, *17*, 773.

van Damme, Eric and Sjaak Hurkens, “Commitment Robust Equilibria and Endogenous Timing,” *Games and Economic Behavior*, August 1996, *15* (2), 290–311.

Vives, Xavier, “Nash equilibrium with strategic complementarities,” *Journal of Mathematical Economics*, 1990, *19* (3), 305–321.

Wilson, Robert B, *Nonlinear pricing*, Oxford University Press on Demand, 1993.

8 Appendix: Bertrand duopoly with differentiated products

To illustrate the idea of a commitment game, consider the differentiated Bertrand duopoly model with linear demand functions. Suppose two symmetric firms ($i = 1,2$) produce differentiated products. The demand for each product depends both on its price and the price of the competitor in a linear way:

$$\begin{aligned} q_1(p_1, p_2) &= 1 - p_1 + \alpha p_2, \\ q_2(p_1, p_2) &= 1 - p_2 + \alpha p_1, \end{aligned}$$

where $\alpha \in (0,1)$ is the parameter that characterizes the degree of products' substitutability: the higher is α, the more substitutable are the products.

If firms set their prices simultaneously and independently, then there exists a unique Nash equilibrium in which both firms charge the following price:

$$p_1^{NE} = p_2^{NE} = \frac{1}{2-\alpha}.$$

As a result, both firms receive profits equal to $\pi_1^{NE} = \pi_2^{NE} = \frac{1}{(2-\alpha)^2}$.

It is easy to see that this pair of profits does not lie on the Pareto frontier. In other words, if both firms decide to increase their price by a small amount ($\varepsilon > 0$), then both profits will increase:

$$\pi_1 = \pi_2 = \left(1 - (1-\alpha)\left(\frac{1}{2-\alpha} + \varepsilon\right)\right)\left(\frac{1}{2-\alpha} + \varepsilon\right) = \frac{1}{(2-\alpha)^2} + \frac{\alpha}{2-\alpha}\varepsilon - (1-\alpha)\varepsilon^2 > \frac{1}{(2-\alpha)^2}$$

for small $\varepsilon > 0$.

Thus, if the firms could agree to coordinate their actions, they could do better. Yet, it is clear that neither firm unilaterally has an incentive to declare it will charge the profit maximizing price $p_1^M = p_2^M = \frac{1}{2(1-\alpha)}$. To see why this pair of prices cannot be a Nash equilibrium, consider the figure.

The figure depicts the isoprofit curves for each firm corresponding to the prices that maximize the firms' joint profit. These isoprofits touch each other at $\left(\frac{1}{2(1-\alpha)}, \frac{1}{2(1-\alpha)}\right)$, the point that maximizes the firms' joint profit. Suppose firm 1 knows that firm 2 will charge p_2^M. If firm 1 charges p_1^M, then it gets $\pi_1^M = \frac{1}{4(1-\alpha)}$. However, if firm 1 charges a different price, in particular, any price in the range $\left(p_1^P, p_1^M\right)$, then its profit will be strictly higher than π_1^M. In other words, firm 1 has a temptation to deviate from the price it is supposed

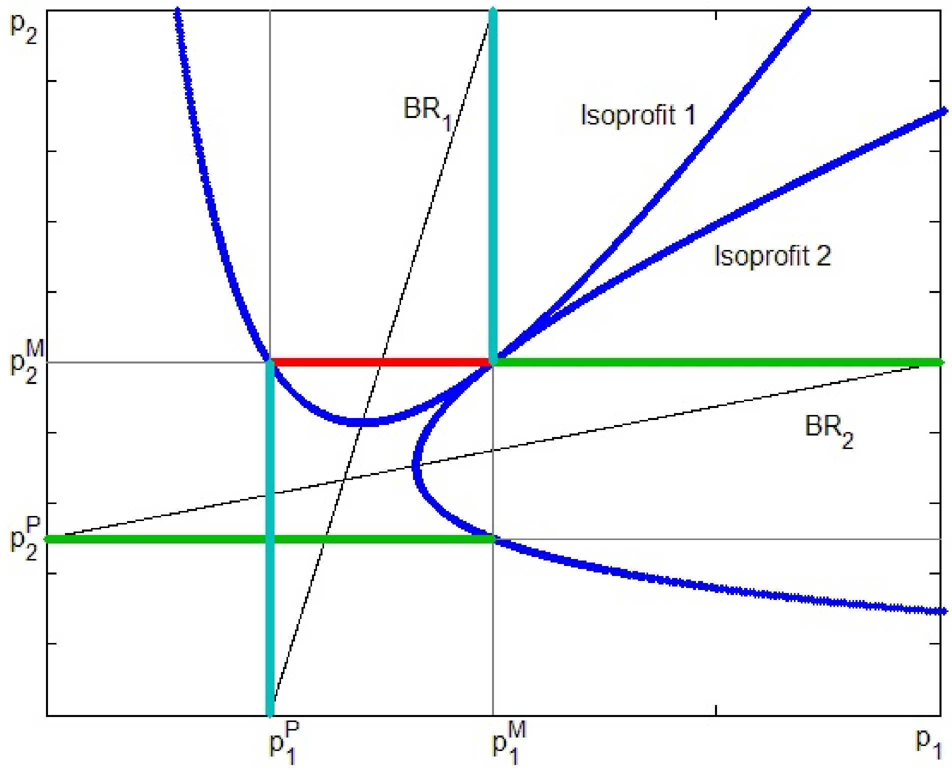

Figure 4: Differentiated Bertrand Duopoly. How to Support a Joint-Profit Maximizing Outcome

to charge and charge something lower. Unless firm 1 can commit not to charge prices from the interval $\left(p_1^P, p_1^M\right)$, the pair of prices $\left(p_1^M, p_2^M\right)$ cannot be an equilibrium.

Suppose now that firms have ability to restrain themselves independently of each other. To be more precise, suppose that firms choose a subset of their prices first. We saw that firm 1 can profitably deviate if it charges any price from $\left(p_1^P, p_1^M\right)$. Therefore, in an equilibrium it needs to restrain itself and exclude all prices from that interval. However, it has an incentive to cheat (a "temptation"). Therefore firm 2 has to be able to punish firm 1 if it includes prices from that interval. Hence, firm 2 has to choose not only price p_2^M that it is supposed to play in an equilibrium but also some other prices that it will use as punishments should firm 1 not commit to exclude prices in the interval $\left(p_1^P, p_1^M\right)$.

For reasons discussed in the main text, it is sufficient in this case for each firm to choose just two prices: the pair of prices that will be played along the equilibrium path $\left(p_1^M, p_2^M\right)$ and two prices that will be used as punishments, namely p_1^P and p_2^P. To see why committing to just two prices in an initial stage will work, suppose firm 1 decides to deviate and restrains itself to some other price. It is easy to verify that firm 2 will charge p_2^M if firm 1 charges any price higher than p_1^M. If firm 1 charges any price lower than p_1^M, then firm 2 should choose p_2^P. Thus, firm 1 cannot benefit from a deviation given firm 2 commits to $\left\{p_2^M, p_2^P\right\}$ since the isoprofit of firm 1 lies to the north of firm 2's optimal response. Perhaps, however, firm 1 could benefit from choosing more than one price at the commitment stage? It turns out that no matter what subset of prices the firm chooses, there will always exist a pure-strategy Nash equilibrium in the corresponding subgame in which firm 2 will play either p_2^M or p_2^P. In neither case, as we already have seen, can firm 1 benefit if it deviates.

Thus, in this widely studied game, if firms can restrain themselves independently of each other, they can coordinate on the outcome that maximizes their joint profit.