Essays in Applied Microeconomics

Inaugural-Dissertation
zur Erlangung des Grades eines Doktors
der Wirtschafts- und Gesellschaftswissenschaften
durch die
Rechts- und Staatswissenschaftliche Fakultät
der Rheinischen Friedrich-Wilhelms-Universität
Bonn

vorgelegt von
SIMON DATO
aus Much

Bonn 2015

Dekan:	Prof. Dr. Rainer Hüttemann
Erstreferent:	Prof. Dr. Matthias Kräkel
Zweitreferent:	Prof. Dr. Jörg Bude

Tag der mündlichen Prüfung: 6.7.2015

Acknowledgements

Over the last years, I have received a lot of support and assistance from various people which helped to improve this thesis significantly. I'd like to take this opportunity to express my gratitude and appreciation. First, and foremost, I would like to thank my first supervisor Matthias Kräkel. Beside placing your trust in me and giving me the opportunity to write this thesis, I am truly happy that you gave me plenty of scope in choosing the topics to work on. Moreover I am highly indebted to your support and constructive criticism leading to a permanent improvement of all chapters contained in this thesis. A further thank you goes to my second supervisor Jörg Budde for taking your time and putting in all the effort to comment on my thesis. To my co-authors Andreas Grunewald, Carsten Dahremöller, Daniel Müller and Petra Nieken: It was and still is a great pleasure to work with you and I can truly say that I have learned a lot from everyone of you. Carsten, you have guided me through writing my first paper which was an enormous help. Petra, you have introduced me to the field of experimental economics and guided me in how to design conduct an experiment. I have really learned a lot from you. Andreas and Daniel, I feel blessed and consider myself lucky having the opportunity to work with you. Doing research with you or talking about economics as well as any other stuff is a great source of inspiration and good fun. Finally, I am outstandingly happy to have great family: I owe a lot of thanks to my parents and my sisters for their unconditional love and support. More than anything, I want to thank Lena for supporting me over the last years and for being the wonderful person that you are!

Contents

List of Figures

List of Tables

Introduction

Quite a few years ago, as a student of economics, several great lectures have drawn my interest for microeconomic theory. I was fascinated by the idea of creating an image of reality, impose simplifying assumptions to ensure tractability and then answer economically relevant questions within a theoretical framework. At the same time, however, I found myself unsatisfied with the following statement that was frequently made in these lectures: "Regarding the economic agent, we impose the assumption [...], which seems somehow unrealistic, but is necessary to be able to solve the model." Among others, the key assumptions imposed on the economic agent are that he is rational, purely selfish and, hence, motivated solely by expected material utility maximization. A large body of empirical evidence exists documenting observed human behavior that is hard to reconcile with the assumptions stated above.[1] I began to worry about the missing congruence with reality of these assumptions and, more importantly, the corresponding theories' low explanatory power and predictive inaccuracy. At this point behavioral and experimental economics come into play. Experimental economics pursues the goals of (i) achieving a better understanding of the foundations of human decision-making, and (ii) testing the validity of economic theories. Behavioral Economics aims at increasing the explanatory power of economics by augmenting the standard model with insights from psychological and experimental research. Hence, the empirical validity of microeconomic theories is tested in controlled laboratory experiments. Behavioral economics, in turn, incorporates insights from experiments into microeconomic theories to increase the degree of realism in these models. The underlying motive is that, ceteris paribus, the higher the degree of realism, the more accurate theoretical insights should be, leading to better predictions and better policy recommendations. It becomes apparent that behavioral and experimental economics are not two distinct areas of research but are very closely related and complement each other. This dissertation employs the methods from experimental and behavioral economics to tackle the following questions in four

[1]For a good survey on the empirical evidence see, among others, DellaVigna (2009).

self-contained chapters: First, how does consumers' limited attention regarding add-on prices affect firms' ability to engage in collusive behavior? Do firms have an incentive to educate consumers and bring add-on prices into focus? Second, do females and males systematically differ in their willingness to sabotage the opponent in a rank-order tournament? Do sabotage choice and performance in a tournament differ with respect to the gender of the opponent? Third, in very general settings of strategic interaction as well as in a specific application, how is the behavior of players and equilibrium outcome shaped by expectation-based loss aversion?

Chapter 1 is joint work with Carsten Dahremöller. We analyze the impact of consumer myopia on competition and firm behavior. In our model, firms repeatedly sell a primary good and a respective add-on. The population of consumers contains a fraction of myopic consumers that do not take into account the price of the add-on when deciding whether or not to purchase the good. Hence, myopic consumers underestimate the total price of the bundle and are more likely to buy the good compared to sophisticated consumers. Every firm may educate myopic consumers and turn them into sophisticated consumers by unshrouding the add-on. We study what impact consumer myopia in the add-on market has on pricing behavior and on the ability of firms to engage in collusion. Our main result is that the existence of myopic consumers facilitates collusion. This result is driven by the finding that a deviation from collusive play is less rewarding if some consumers are myopic. To provide an intuition for this result, recall that, in the absence of consumer myopia, a deviating firm is able to attract all consumers and obtains the whole collusive profit by slightly undercutting the collusive price. If some consumers are myopic, however, first note that colluding firms will sell the base-good for the lowest possible price in order to trick myopic consumers into buying. A deviating firm is then able to undercut only in the add-on dimension and needs to decide whether or not to unshroud the add-on. If she decides not to unshroud the add-on, she attracts only the sophisticated consumers since myopes do not recognize the prize cut. By unshrouding, she educates all myopics to sophisticated consumers such that all consumers recognize the price cut of the deviating firm. At the same time, however, unshrouding crowds formerly myopic consumers out of the market since they realize that the total price is higher than anticipated and therefore abstain from buying. Either way, a deviating firm is not able to obtain the whole collusive profit. Therefore, consumer myopia makes a deviation from collusive play is less rewarding and, hence, facilitates collusion. Moreover, consumer myopia leads to higher collusive profits. Firms therefore have a strong incentive to leave the add-on market shrouded: Shrouding might be a requirement for collusion in the first place. Moreover, even if collusion is already stable, a shrouded market leads to higher collusive as well as individual profits. These results yield important implications for governmental policies and regulatory interventions. First, a shrouded market in which no firm educates consumers is a sign for cartelization. Hence, the observed obfuscation in a market can serve as a proxy signal for illegal industry agreements. Second, a regulatory intervention with the aim to unshroud the add-on is increasing total welfare if it can lower the fraction of myopic consumers sufficiently.

Chapter 2 is joint work with Petra Nieken. We study the differences in behavior of males and females in a two-player tournament with sabotage in a controlled lab experiment. Each contestant produces output by carrying out a real-effort task. In addition, each contestant is given the opportunity to engage in a costly sabotage activity. Sabotage partially destroys the opponent's output and, hence, increases the own winning probability. The contestant with the higher final output then receives the winner prize. We additionally implemented a principal who is paid based on the agent's output. Hence, sabotaging the opponent imposes a negative externality on the principal's payoff. We find that males and females do not differ in their achievements in the real effort task but in their choice of sabotage. Regarding sabotage decisions, however, we find a strong gender gap. On average, males select twice a much sabotage than females. This gender gap in sabotage is highly significant. Given that females and males perform on the same level, the gender gap implies that males win the tournament more often than females. It turns out that profits of females and males do not differ significantly from each other since sabotage is costly. The same gender gap that we find in actual sabotage choices is as well present in the beliefs about the opponent's sabotage decision. Males not only sabotage their opponent more strongly, they also expect to be sabotaged by the opponent to a higher degree than females do. We conducted a control treatment where contestants were able to cheat by adding extra points on their own output instead of destroying the opponent's output via sabotage. By replacing the sabotage activity with a cheating activity, the positive effect on the own winning probability remains unchanged, but the negative externality on the principal's payoff is turned into a positive externality. Nevertheless, males again choose significantly higher cheating levels than females, which suggests that the gender gap in sabotage is not driven by differences in social preferences towards the principal. In the gender treatment, we revealed the gender of the opponent before the tournament starts. The gender gap in sabotage is persistent and, in addition, males now perform significantly better than females. Compared to the other treatments, females perform on the same level but males increase their performance compared to the other treatments. This effect is especially pronounced if the opponent is female. We discuss possible explanations for our findings and their implications.

Both the third and the fourth chapter are joint work with Andreas Grunewald and Daniel Müller. Both papers analyze the behavior of expectation-based loss averse individuals in settings of strategic interaction. Next to Expected Utility Theory, Prospect Theory with the core features of reference dependence and loss aversion has become the most widely applied approach for modeling risk preferences. One drawback of Prospect Theory is that it remains completely agnostic about how the reference point is formed. Kőszegi and Rabin (2006, 2007) fill this gap by proposing a framework of how the reference point is endogenously determined by rational expectations. They define two different solution concepts, personal equilibrium (PE) and choice-acclimating personal equilibrium (CPE). PE is the appropriate solution concept in a situation where the decision maker has formed fixed expectations about what to do some time before actually reaching a decision, whereas CPE applies to a situation in which the decision maker is confronted with the decision rather unexpectedly.

Chapter 3 provides a comprehensive analysis regarding strategic interaction under expectation-based loss-aversion. First, we develop a coherent framework for the analysis by extending the equilibrium concepts of Kőszegi and Rabin (2006, 2007) to strategic interaction and demonstrate how to derive equilibria. Second, we delineate how expectation-based loss-averse players differ in their strategic behavior from their counterparts with standard expected-utility preferences. In particular, we derive three behavioral features of expectation-based loss averse players which qualitatively differ from the strategic behavior of players with standard preferences. Recall that a player with standard expected utility preferences is willing to play any mixture over a set of pure strategies if she is indifferent between these pure strategies. After a slight change in the opponents' strategies, however, she will not be indifferent between these strategies which completely wipes out the willingness to play a mixture over this set of pure strategies. For fixed expectations, expectation-based loss averse players are *adaptive* in the sense that they are potentially willing to play a mixture over a set of pure strategies for a nontrivial range of opponents' strategies. Moreover, expectation-based loss averse players exhibit the behavioral feature of *decisiveness*, according to which, for given strategies of the other players, they are willing to play at most one mixture over a given set of strategies. Finally, under choice-acclimating expectations, players are *reluctant* to play a mixed strategy irrespective of the game and the other players' strategies. Third, we analyze equilibrium play under expectation-based loss aversion and comment on the existence of equilibria.

In chapter 4, we apply the solution concept defined in the previous chapter to rank-order tournaments and comment on the work of Gill and Stone (2010). Many insights regarding rank-order tournaments rest upon contestants' behavior in symmetric equilibria. As shown by Gill and Stone (2010), however, a symmetric equilibrium may not exist if contestants are expectation-based loss averse. We complement this important finding by delineating the circumstances under which a focus on symmetric equilibria is nevertheless justifiable. First, the existence of a symmetric equilibrium is guaranteed if the contestants' concern for psychological utility does not outweigh their concern for material utility and minimal effort comes without costs. Second, if contestants enter the tournament with lagged fixed expectations rather than choice-acclimating expectations, symmetric equilibria exist irrespective of the contestants' degree of loss aversion.

1. Collusive Shrouding and Cartelization

1.1. Introduction

In many markets product information is not easily available and consumers have difficulties when trying to collect the information that is relevant for their shopping decisions. In the terminology of the literature, building on the seminal paper of Gabaix and Laibson (2006), these are termed shrouded markets. Firms that participate in these markets can either foster or alleviate the degree of obfuscation. In the economic literature there is ongoing debate about what firms should do in such a situation. Gabaix and Laibson (2006) argue in a model of add-on markets that firms will not reveal add-on prices if there are enough myopic consumers, i.e. if there are enough consumers that do not incorporate all information that is available to them. If any firm unshrouds, some of the myopic consumers get educated and add-on prices get revealed. However, it turns out that the education of consumers does not have any strategic effect in their model since it happens so late that it does not have any impact on the game and on the incentives of the competing firms. Dahremöller (2013) picks up this point and shows that, if the education of consumers has strategic implications for the game, the result that firms have an incentive to shroud the add-on breaks down.

We further expand this framework by designing an infinitely repeated game in which consumer education is a strategic variable in the sense that it has an effect on the payoffs and strategies of firms. We show that the natural equilibrium[1] of the game is one in which firms set competitive prices and unshroud the add-on. However, for sufficiently high discount factors there also exists an equilibrium in which firms collude on monopoly pricing.[2] In particular, if all consumers are sophisticated and if the discount factor is high enough, firms can collude on monopoly pricing.

[1]The term 'natural equilibrium' refers to an equilibrium in which firms play the equilibrium strategy of a respective one-shot game in each period of the game.

[2]As Fudenberg and Maskin (1986) and many others have shown, infinitely repeated games give rise to a plethora of equilibria besides the natural equilibrium.

One of our central findings is that the existence of myopic consumers makes the collusive equilibrium more stable in the sense of lowering the critical discount factor for which collusion is sustainable. In terms of the model this implies that, dependent on the market constellation, firms have an incentive not to unshroud the add-on in order to keep the fraction of myopic consumers high enough for collusion to be stable. In addition, even if collusion is already stable with only sophisticated consumers, shrouding increases monopoly profits.

The fact that shrouding makes collusion more stable and, in many cases, is even a prerequisite for collusion has strong implications for competition analysis and antitrust regulations. First, since, dependent on parameter constellations, shrouding is a requirement for collusion, the regulator might intervene in order to decrease market obfuscation in order to destabilize collusion. Second, since again shrouding might be a requirement for collusion, the degree of obfuscation in a market might be a proxy signal for ongoing collusion in the market. Hence, since it is usually difficult to detect agreements on collusive pricing, it might be helpful for the regulator to consider the degree of obfuscation as an additional indicator for collusion.

One example of firms coordinating on intransparency with regard to their products was the Lombard Club, which was a cartel of Austrian banks that was detected by the European Commision (see European Commission (2004)). For example, at a meeting of the involved banks in 1994 “everyone agreed that, if questioned by the press or by the Association for Consumer Information for rate comparison purposes, they should in future stick to communicating only the (official) rates posted at the counter and not answer any further questions.” In another agreement in 1996, the involved banks coordinated on valuing and pricing their portfolio lists only in Austrian Schilling while dropping any reference to the Euro. It is documented that the involved banks agreed that valuing and pricing in both Euro and Austrian Schilling would be more transparent to consumers, but they deemed that competition in this dimension should be avoided. In 1999 the involved banks agreed not to publisize a comparison of their savings products since this would open a way to “fresh competition.”

More evidence for the connection between shrouding and collusive profits is presented by Brandenburger and Nalebuff (1998) for the U.S. airlines industry. In 1992, from an initial situation in which all pricing schemes in the industry were rather opaque and intransparent, American Airlines started a new pricing scheme “which emphasizes simplicity and equity and value.” Its competitor United Airlines responded within forty-eight hours and in the following all other major competitors also quickly adopted simplified pricing schemes. Just three days later, another competitor, Trans World Airlines, revised its pricing schedule and severely undercut industry prices, which was again followed by quick price cuts of all major competitors. This example indicates that there might be a tight relationship between the increase of transparency by the first firm and the following cascade of increased transparency and price cuts by the other firms in the market. This linkage between obfuscation and pricing will also be present in our results.

In our model, we analyze markets for goods whose total price consists of more than one element, for example markets for a base good and an add-on. If both the base good and the add-on are consumed, the effective total price of the product bundle is the sum of both prices. For instance, if a consumer considers buying a printer, she will not only have to pay the immediate price for the printer, but most likely have future expenses for compatible refill cartridges. Another application for our framework are goods that trigger future payments, for example subscriptions for which the total price is the (discounted) sum of all payments.

If a consumer wants to correctly calculate the effective total price of the product bundle, she needs to possess all relevant information and therefore a high degree of consumer attention. If the attention of the consumer is limited though, she possibly does not fully recognize the effective total price. Consumers that exhibit this bias are called myopic consumers. As a result of their myopia, they may not be able to make rational consumption decisions. For example, Cruickshank (2000) reports that users of a current account seldomly fully understand all details of the contract and in most cases pay only little attention to add-on fees or other contract specifications. In line with these findings are the results from a survey considering consumer empowerment in the European Union which was conducted by TNS (2011). Addressing the question whether European citizens are sufficiently empowered as consumers, it is reported that almost six out of ten interviewees did not fully read the terms and conditions of the latest service contract that they signed (including contracts for gas, electricity, mobile phones, bank accounts, or insurances). Building on these results we assume that the consumer population contains a positive fraction of myopic consumers.

Considering such markets, it seems reasonable to assume that firms can exert some influence on the degree of obfuscation. Note here that in many markets there is little scope for obfuscation in the base good market since transparency in this dimension is necessary to attract consumers to the market. This is different for add-on markets which usually offer greater possibilities for firms to shroud information and prices. One example for such a market would be monthly subscriptions. For example, firms can either employ a transparent flat fee or they can employ a payment structure that is increasing over time without prominently mentioning it in their marketing material, which may mislead consumers to anticipate a lower price than they end up paying. In terms of the model this means that firms can either shroud or unshroud the add-on. If all firms shroud the add-on, a fraction of $\alpha \in (0, 1]$ consumers is myopic and does not consider the add-on price. If any firm unshrouds the add-on, all myopic consumers get educated and behave like the sophisticated consumers for the current and all future rounds of the game.[3] In

[3] Note here that we abstract from the possibility that firms can actively increase the fraction of myopic consumers α. In reality it seems reasonable that firms can somehow increase α. While we do not explicitly model this possibility, it will turn out that if a shrouding equilibrium exists, firms will have an incentive to increase α.

contrast to Gabaix and Laibson (2006) and Dahremöller (2013) we assume that the total demand for the product bundle is not fixed. Instead, consumers have a personal valuation for the product bundle, which is heterogeneous over the consumer population. If the valuation of a consumer is lower than the anticipated price of the product bundle, she will prefer not to buy. Note here that for given prices the total demand is higher if there are myopic consumers than if there are only sophisticated consumers. This is because myopic consumers underestimate the total price of the product bundle and hence are more likely to participate in the market.

Our main finding is that the existence of myopic consumers facilitates collusion. In terms of the model this implies that the critical discount factor for which collusion is stable is decreasing in the fraction of myopic consumers. The main driver of this finding is the fact that a deviation from collusion is less rewarding if many consumers are myopic. To get an intuition for this result consider a situation in which firms collude on monopoly pricing. In a model with only sophisticated consumers a firm that deviates and undercuts the collusive price attracts all consumers in the market and earns the entire monopoly profits. This is different if some consumers are myopic. Suppose first that a firm deviates by only lowering its price. Then myopic consumers do not perceive the change in the price and therefore will not switch to the deviating firm. Suppose second that a firm deviates by lowering its price and by unshrouding the add-on. Such a deviation would attract all consumers that still participate in the market, but would lower the total demand since some myopes realize that prices are higher than they anticipated. Both these effects make a deviation from collusion less rewarding and hence collusion is more stable. In other words, the existence of myopic consumers facilitates collusion.

In addition, even if collusion is already stable for a given population composition, firms have an incentive to continue increasing the fraction of myopic consumers. Since myopic consumers underestimate the total price of the product bundle, shrouding may trick these consumers into consumption. Hence, the total demand for the product bundle and the profits of the firms are rising with the fraction of myopic consumers.

Our analysis also yield several insights on welfare. If the consumers valuation for the product bundle is too low, a decision to buy the product bundle is inefficient and will be regretted by the consumer ex post. Therefore we find that shrouding has a negative impact on consumer welfare. This result has implications when applying our results to a regulatory perspective. We find that a regulatory intervention with the aim to unshroud the add-on is increasing total welfare if it can lower the fraction of myopic consumers α sufficiently. This finding is in line with Kosfeld and Schüwer (2011) who analyze the framework of Gabaix and Laibson (2006) from a regulatory perspective and show that in their framework a regulatory consumer education can have positive as well as negative effects on welfare. In addition, we find that a regulation to unshroud the add-on is always increasing the consumer surplus.

When considered from a regulatory perspective, our results suggest that regulatory tools with the aim to unshroud the add-on can also impede collusion. If a regulatory intervention can decrease the fraction of myopic consumers α sufficiently, collusion can potentially be prevented or destabilized. Examples for such intervention would be informational campaigns to increase consumer sophistication or regulations that enhance market transparency. Even if such efforts to unshroud cannot decrease the fraction of myopic consumers α sufficiently in order to prevent collusion, they still increase consumer welfare as they either prevent consumers from making irrational choices or decrease the prices that consumers pay for the product bundle.

In addition to suggesting tools to impede collusion, our results also suggest new tools to detect collusion. We find that in many parameter constellations shrouding is necessary for collusion to be stable. Hence, a shrouded market is a potential indicator for illegal industry agreements. These markets then are candidates for increased scrutiny and inspections by governmental trustbusters. Traditional antitrust provisions like unannounced inspections or leniency policies were used in order to detect, prove, and prevent collusive industry behavior. However, historic evidence suggests that these tools were only partly successful in preventing collusion and cartelization. In particular, unannounced inspections and leniency policies were mostly targeted at disintegrating existing cartels. Our results suggest a new approach to detect active and intact cartels and prevent future cartel formation.

The remainder of this paper will proceed as follows. In section 1.2 we give a short overview over the related literature. In section 1.3 we present the main analysis and results. Section 1.4 will conclude.

1.2. Review of the Literature

The economic discourse on information disclosure by competitive firms was started by Grossman (1981) and Milgrom (1981). In their works rational consumers are imperfectly informed about product attributes and firms can credibly reveal the missing information. Within this framework, the authors show that competitive firms indeed have an incentive to reveal the missing information since this has a positive effect on their demands and profits.

One of the first works on obfuscation in add-on markets is provided by Shapiro (1995). He argues that there does not exist an equilibrium in which firms shroud the add-on. Shapiro argues verbally that there is a customer winning effect of unshrouding, which implies that a shrouding equilibrium is not stable.

Gabaix and Laibson (2006) were the first to question that view. Building on the work of Shapiro (1995), they use a model of add-on markets and assume that a given fraction of myopic consumers does not consider the add-on price in their purchase decisions. However, each firm can educate the myopic consumers and thereby help them to make more sophisticated decisions. The authors show that if the fraction of myopic consumers is large enough, there exists an equilibrium in which no firm has an incentive to educate consumers and all firms shroud the add-on.

This point is picked up by Dahremöller (2013) who shows that the results of Gabaix and Laibson (2006) are strongly based on their modeling of unshrouding. In particular, Gabaix and Laibson (2006) use a single-period model in which firms can unshroud only in the last stage of the game. This implies that if any firm unshrouds the add-on, the other firms cannot react to this deviation. Also, unshrouding has no impact on any future payoffs since the game ends at this point. In terms of the model this implies that the education of consumers has no consequence for the play of the game and for the strategies of firms.[4] Examining these effects, Dahremöller (2013) shows that if the education of consumers is modeled to have strategic implications for the game, a shrouding equilibrium no longer exists.

If one would transfer the framework of Gabaix and Laibson (2006) to a repeated game, unshrouding would have only short-term effects, while any long-term effects would be neglected. However, in our model we will follow Dahremöller (2013) by assuming that unshrouding has long-term implications. Examples for long-term effects of consumer education include that firms condition their behavior on the play of the previous rounds and include that unshrouding permanently alters the consumer structure in future periods.

1.3. The Model

1.3.1. Model Setup

We model an infinitely repeated game. In each period $n \geq 3$ symmetric firms produce a base good and an add-on at zero costs. Each firm i sets a base-good price p_i and an add-on price $\hat{p}_i$. The common discount factor of firms is $\delta \in [0, 1]$. The consumer population has mass 1 and consists of $\alpha \in (0, 1]$ myopic consumers and $1 - \alpha$ sophisticated consumers. The fraction of myopic consumers only considers the base good prices p_i and neglects the add-on prices $\hat{p}_i$. The remaining fraction of sophisticated consumers is fully informed, rational, and considers both the base good prices p_i and the add-on prices $\hat{p}_i$. We assume that there exists a maximum price $\bar{p}$ for the add-on which can be interpreted as the cost of a last minute substitution or a regulatory usury ceiling.

As outlined before, we assume that in each period each firm can unshroud the add-on. If one firm does so, the myopic consumers get educated, which means that they behave like sophisticated consumers for the current round and all remaining rounds of the game. Consumers derive

[4] In essence, Gabaix and Laibson (2006) show that firms potentially do not have an incentive to reveal their add-on prices. This result holds irrespective of whether this revelation of add-on prices is linked with an education of the consumer population.

utility v from consuming one unit of the base good and the respective add-on in a given period. If a consumer abstains from buying the product bundle, she gets zero utility. Each consumer buys at most one unit of the base good and one unit of the add-on. The 'realized' utility of consumers is $U = v - p_i - \hat{p}_i$. However, myopic consumers mistakenly anticipate to get a utility from buying at firm i of $U^m = v - p_i$. Hence, the myopic consumers do not anticipate their future need for the add-on or, equivalently, anticipate that the add-on price is zero. In contrast, sophisticated consumers correctly anticipate that their utility from buying at firm i is $U^s = v - p_i - \hat{p}_i$. The consumption utility v of each consumer is stochastic with $v \sim U[0, \overline{v}]$. The cdf of v then takes the following form:

$$F(z) = \begin{cases} 0 & \text{if } z \leq 0, \\ \frac{z}{\overline{v}} & \text{if } 0 < z < \overline{v}, \\ 1 & \text{if } \overline{v} \leq z. \end{cases}$$

Since firms compete via Bertrand competition, a consumer buys the bundle at the firm that yields her the highest anticipated utility. If there are several firms that yield the highest anticipated utility, the consumer will choose any one of them with equal probability. In addition, if no firm yields positive anticipated utility, the consumer abstains from buying. The well known result of a one-period game with Bertrand competition is that firms earn zero profits. However, the infinite repetition of a game usually allows for a plethora of equilibria and firm strategies[5] and there exists no general mechanism for equilibrium selection. However, in the following, we assume that if firms collude on pricing or shrouding, they will coordinate on the equilibrium that yields the highest profit per firm.[6] In particular, since firms are symmetric, we focus on equilibria that yield the highest aggregate firm profit.

One implication of the existence of myopic consumers is that firms have an incentive to set high add-on prices along with low base good prices. Now recall the utility functions U^m and U^s. For given prices p_i and $\hat{p}_i$, a firm can always increase the attractiveness of its product bundle for myopics while leaving the attractiveness of the bundle for sophisticates unchanged. The firm can achieve this by lowering the base good price p_i by a small amount and increasing the add-on price $\hat{p}_i$ by the same small amount.

However, this logic of lowering base good prices and raising add-on prices is potentially limited. The reason is that there are several arguments that lead to a lower bound for base good prices. For example, in real-world markets, prices cannot be negative due to possible arbitrage opportunities. If consumers receive money for the purchase of a good, they will buy as many units of the good as they can, creating unlimited profits for themselves and a loss

[5]See, for example, Fudenberg and Maskin (1986).

[6]We will later argue that an equilibrium with less-than-optimal profits would not be more stable in terms of collusion. Hence, there is no obvious reason for firms to collude on less-than-optimal profits.

for the firm that sells the good. In addition to the condition that prices must be non-negative, there potentially also exist reasons for positive lower price limits. One reason for a lower price bound is the possibility to resell parts of the base good. For instance, if consumers could buy a printer at zero costs and sell the copper wires or other parts of the printer for a profit, they might exploit this as an arbitrage opportunity. Another argument for price boundaries in the base good dimension is brought forward by Miao (2010) who argues that if the base good and the add-on are substitutes, there will be a lower limit for the base good price. Printers, for example, are sold with a starting cartridge. If a cartridge runs low, the consumer has the choice between buying a new cartridge or buying a new printer that is already equipped with a new starting cartridge. If the printer is very cheap compared to the refill cartridge, firms will not be able to sell their high priced refills. Miao (2010) shows that this creates a lower limit for the base good price.

Following the above argumentation, we impose a lower bound for the base good price. For simplicity we set this limit to 0. We will later show that this lower bound for the base good is reached if the following condition holds:

Assumption 1.1. $\overline{v} \leq \bar{p}$.

We will assume this condition to hold for the remainder of the paper. Note that the assumption of a lower bound for the base good price creates results that stand in contrast to the traditional Chicago school argument on add-on pricing. Formalized, for example, by Lal and Matutes (1994) and Gabaix and Laibson (2006), the Chicago school argument reckons that high profits in the add-on dimension are fully competed away in the base good dimension. So suppose that, in a market with only a base good, firms would charge an equilibrium price of $p^{\dagger}$. Now suppose that an add-on is introduced, yielding an equilibrium add-on price of $\hat{p}^{\dagger}$. Then the Chicago school argument predicts that the new base good price will simply be the old base good price minus the new add-on price, i.e. $p^{\dagger} - \hat{p}^{\dagger}$. In other words, the Chicago school argument predicts that the base good fully subsidizes the add-on. Obviously, such a cross-product subsidization is not always possible if a lower boundary for the base good price exists.

1.3.2. Analysis

To determine the effect of the existence of myopic consumers we now want to compare a situation in which all consumers are sophisticated to a situation in which a fraction of $\alpha > 0$ consumers is myopic. We will then show that equilibria exist in which firms collude on shrouding. In particular, there exist constellations in which shrouding is necessary to allow firms to cartelize and thereby to jointly earn monopoly profits.[7]

[7]Note here that the main focus of our model is to test whether market obfuscation facilitates collusion. It may be the case that there exist other equilibria, which are not considered in our model, in which firms do not coordinate on prices but still shroud the add-on.

Only sophisticated consumers

First suppose that all consumers are sophisticated. Sophisticates take both the price for the base good and the price for the add-on into account. We now want to determine for which parameter constellations firms can cartelize. If firms can coordinate on monopoly pricing, they will set monopoly prices p^M and $\hat{p}^M$ and earn aggregate monopoly profits π^M. The profits π^M are split up equally between the firms such that each firm earns a profit of $\frac{1}{n}\pi^M$. If a firm expects to earn $\frac{1}{n}\pi^M$ for all future periods, the present value of these cash flows is $\frac{\pi^M}{n(1-\delta)}$. If any firm deviates from monopoly behavior and undercuts marginally, it will attract all consumers and make a deviation profit of π^{dev}. Since the deviating firm can undercut the monopoly prices only marginally and thereby attract all consumers, it would earn a deviation profit of the entire monopoly profits $\pi^{dev} = \pi^M$. This deviation will trigger a grim-trigger punishment by the other firms.[8] Hence, after such a deviation collusion breaks down and from that point onwards firms will compete via Bertrand competition and make zero equilibrium profits $\pi^{NC} = 0$ for all following periods (NC=non-collusive). The present value of the cash flows after deviation hence is $\pi^{dev} + \frac{\delta\pi^{NC}}{1-\delta}$. Now we want to determine the critical discount factor, which is the discount factor for which firms are indifferent between sticking to collusion and deviating from collusion. Applying our results, the critical discount factor is given by the solution of the following equation:

$$\begin{aligned} \frac{\frac{\pi^M}{n}}{1-\delta} &= \pi^{dev} + \frac{\delta\pi^{NC}}{1-\delta} \qquad (1.1) \\ \Rightarrow \delta^* \quad &= \frac{n - \frac{\pi^M}{\pi^{dev}}}{n} = \frac{n-1}{n}. \end{aligned}$$

Thus, for all discount factors $\delta \geq \delta^*$ collusion is sustainable.[9]

The impact of myopic consumers

Now suppose that a fraction of $\alpha > 0$ consumers is myopic. Recall that, in contrast to sophisticated consumers, myopes do not consider the add-on prices $\hat{p}_i$. If all firms cooperate and charge prices p and $\hat{p}$, their aggregate profit is given by:

$$\pi^M(p, \hat{p}) = [\alpha\,(1 - F(p)) + (1-\alpha)(1 - F(p+\hat{p}))](p+\hat{p}).$$

It follows that if firms collude, they will set the base good price equal to its lower bound:

[8]See, for example, Mailath and Samuelson (2006) for a detailed discussion of grim-trigger strategies.

[9]Note here that a cooperation would not be more stable if firms would coordinate on profits that are lower than the monopoly profits. A deviating firm can always earn the aggregate collusion profit by undercutting marginally. Examining condition (1.1), there exists no coordination on lower-than-optimal profits that makes a deviation less rewarding relative to the collusion profits. Furthermore, coordinating on higher prices than in the monopoly case does not help either since this would destabilize collusion as a deviating firm could undercut and set monopoly prices, thereby earning higher than collusive profits.

Lemma 1.1. *Suppose $\overline{v} \leq \bar{p}$ holds. Then colluding firms will set their base good price at its lower bound, i.e. $p^M = 0$.*

The proof is contained in the Appendix. The intuition for this result is as follows. Suppose that firms set prices such that some sophisticates still participate in the market ($p+\hat{p} < \overline{v}$). Then for every base good price $p > 0$, the firms have an incentive to lower the base good price p and to increase the add-on price $\hat{p}$. This would leave the total bundle price and the demand from sophisticates unchanged, but would increase the demand from myopic consumers. Suppose in contrast that firms set prices such that no sophisticated consumer participates in the market ($p + \hat{p} > \overline{v}$). Then in turns out that the add-on is profitable enough such that firms do not want to decrease demand by increasing the price of the base good.

Applying the result of Lemma 1.1 ($p = 0$) to the aggregate profit function yields the following collusive profit:

$$\pi^M(\hat{p}) = [\alpha + (1-\alpha)(1 - F(\hat{p}))]\hat{p}.$$

In the following analysis we have to distinguish between an 'inner solution' in which both types of consumers buy the product bundle ($\hat{p} < \overline{v},\ F(\hat{p}) < 1$) and a 'corner solution' in which only myopic consumers buy the product bundle ($\hat{p} \geq \overline{v},\ F(\hat{p}) = 1$). If we have an inner solution, both consumer groups have positive demand for the product bundle and the product bundle yields positive utility to some consumers. In contrast, the corner solution is characterized by an add-on price $\hat{p}$ that, if it would be fully considered, exceeds every consumers valuation. In this case only myopics possibly consume the product bundle.

Now we want to derive the global maximum of the profit function. First suppose that firms play an inner solution. Then, the aggregate profit of firms takes the following form:

$$\pi^M(\hat{p}) = \left[\alpha + (1-\alpha)\left(1 - \frac{\hat{p}}{\overline{v}}\right)\right]\hat{p}.$$

The aggregate profit is maximized by charging the monopoly price $\hat{p}^M = \frac{\overline{v}}{2(1-\alpha)}$. Given this price, profits are given by

$$\pi^M_{inner} = \frac{\overline{v}}{4(1-\alpha)}. \tag{1.2}$$

Now consider the corner solution with $\hat{p}^M > \overline{v}$. In this case firms only sell the product bundle to myopic consumers. Then the profit function of firms is given by:

$$\pi^M(\hat{p}) = [\alpha]\,\hat{p}.$$

Obviously, it is optimal to charge the highest possible add-on price $\hat{p}^M = \bar{p}$, yielding a profit of:

$$\pi^M_{corner} = \alpha\bar{p}. \tag{1.3}$$

It will then depend on parameter constellations whether firms will prefer the equilibrium with the inner solution or the equilibrium with the corner solution.

Note here that the optimal price of the inner solution $\hat{p}^M = \frac{\overline{v}}{2(1-\alpha)}$ is larger than the maximum valuation $\overline{v}$ if $\alpha > \frac{1}{2}$. In this case it holds that $1 - F(\hat{p}^M) = 0$ and therefore no sophisticated consumer buys the product bundle. Hence, for $\alpha > \frac{1}{2}$ the inner solution is not feasible in the sense that the derived maximum does not lie in the specified interval. If that is the case, the corner solution will be the global profit maximum.

To get an intuition for the form of the profit function we have depicted two possible constellations in the following graphs: For a sufficiently low share of myopic consumers $\left(\alpha \leq \frac{1}{2}\right)$,

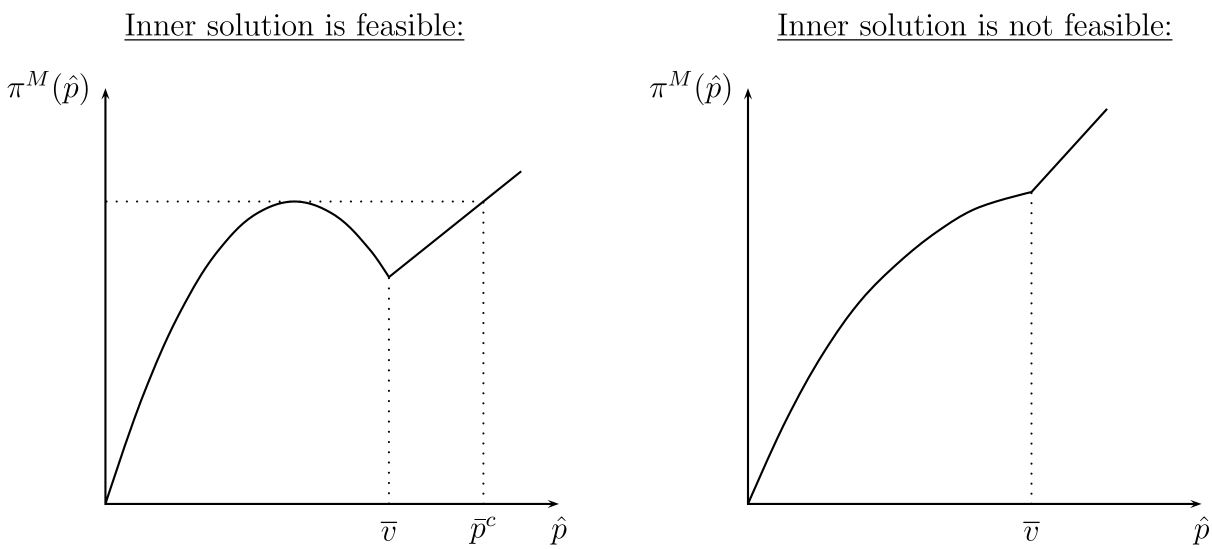

Figure 1.1.: Two possible functional forms of the aggregate firm profits.

the inner solution is feasible in the sense that the profit-maximizing price of the inner solution does not exceed the maximum valuation $\overline{v}$. The inner solution then corresponds to the global maximum of the profit function if the upper bound for the add-on price is not too large, i.e. if $\bar{p} \leq \bar{p}^c$. If, however, the fraction of myopic consumers is large enough, the maximum of the inner solution is not feasible anymore. In this case, only the maximum of the corner solution can be optimal.

When examining the profit functions, we see that the aggregate profit for both the inner solution and the corner solution are increasing in the fraction of myopic consumers α. This leads to the following result:

Proposition 1.1. *Monopoly profits are increasing in the share of myopic consumers α.*

This finding mainly stems from the impact that myopic consumers have on the demand function. Myopic consumers are always more likely to buy the product bundle since they underestimate its total price. Hence, in both the inner and the corner solution, the demand for the product bundle is increasing in the fraction of myopic consumers. This directly implies that the monopoly profits of firms are also increasing in the share of myopic consumers.

If firms cartelize and coordinate on monopoly prices, they are able to maximize aggregate profits, which will then be split up equally among them. Clearly, all firms prefer these monopoly profits over perfect competition with zero equilibrium profits. Nevertheless, there may be an individual short-term incentive to deviate from monopoly pricing: A firm may deviate by either unshrouding the add-on and/or by setting a lower or a higher price than the one that was set in the collusive phase. Lemma 1.1 implies that if a firm wants to deviate from collusion and attract further customers, it can only do so by lowering it's add-on price, but not by lowering its base good price. Note that it is not obvious whether an optimal deviation involves unshrouding the add-on. This is because unshrouding potentially has partly negative effects since, if myopes are turned into sophisticates, they might refrain from buying the product bundle. Therefore, for given prices, unshrouding is decreasing the demand for the product bundle. In the following, we will show that despite its negative effect on demand, an optimal deviation from collusion comprises unshrouding the add-on. The reason is that, due to low base good prices and high add-on prices, firms generate their profit through add-on sales. However, myopic consumers do not incorporate the add-on prices into their purchase decisions. Hence, myopes do not react if the deviating firm changes its add-on price. This creates an incentive to unshroud, since, once myopes are educated, they react to the change in the add-on price. In the following, we will show that this effect dominates the aforementioned reduction in demand and a deviating firm will unshroud the add-on.[10]

First note that firms will never deviate by increasing the add-on price. If firms play a corner solution and charge the highest possible add-on price, they are simply not able to raise the price. If firms play an inner solution and a firm deviates by raising its add-on price, sophisticates will prefer to buy from the other firms and only myopic consumers possibly buy the product bundle from the deviating firm. This is because myopes do not take the add-on price into account and therefore will not change their behavior after a change in the add-on price. Hence, the

[10] Note here that the result that a deviating firm will always unshroud the add-on may be an artefact of our particular population distribution. However, while this result is convenient for the analysis, it is not necessary for our results. To get an intuition for this, recall that if the add-on is unshrouded, firms make zero profits in the competitive equilibrium. Now suppose that firms initially collude and any firm deviates, but does not unshroud the add-on. Then, the other firms could unshroud the add-on as part of their grim-trigger strategies. Since there are at least two firms that do so $(N \geq 3)$, they have no individual incentive not to unshroud the add-on. Hence, if any firm deviates, but does not unshroud, the non-deviating firms will react by unshrouding the add-on and the following non-collusive profits will again be zero. Compared to a situation with only sophisticated consumers, unshrouding still is less attractive with the existence of myopic consumers because a deviating firm does not attract all consumers.

deviating firm will optimally set the maximum add-on price $\hat{p}^{dev} = \bar{p}$, yielding a deviation profit of $\pi^{dev} = \frac{1}{n}\alpha\bar{p}$. It then holds that $\pi^{dev} \leq \frac{1}{n}\pi^M$ and a deviation yields lower profits than the profits earned by sticking to the collusive play.[11] Therefore, increasing the add-on prices is not a profitable deviation.

We can conclude that if a firm decides to deviate, it will undercut the collusive add-on price. Note that it is not possible to undercut the base good price since it is already at its lower bound. Hence, the firm can only undercut in the add-on price dimension. If a firm undercuts the add-on price, it attracts all sophisticated consumers. In particular, the deviant firm has two possibilities to undercut the collusive add-on price:

The first possibility is that the firm undercuts the add-on price and unshrouds the add-on, thereby educating all myopic consumers. This lures a larger share of consumers to the deviant firm because the fraction of price sensitive consumers has increased. At the same time unshrouding potentially crowds many formerly myopes out of the market because, by taking the add-on into account and learning about higher than anticipated add-on prices, some myopes realize that they would receive negative utility from buying and therefore decide to refrain from the market.

The deviation profit that results if a firm unshrouds the add-on takes the following form:[12]

$$\pi_i^{dev} = \left[1 - \frac{\hat{p}_i^{dev}}{\overline{v}}\right] \hat{p}_i^{dev}. \tag{1.4}$$

Maximization of (1.4) yields:

$$\hat{p}_i^{dev} = \frac{\overline{v}}{2}.$$

This price is feasible regardless of the strategies that firms played in the collusive phase ($\hat{p}_i^{dev} \leq \overline{v}$ and $\hat{p}_i^{dev} \leq \hat{p}^M$). Inserting the price into the profit function then yields a deviation profit of

$$\pi_i^{dev} = \frac{\overline{v}}{4}.$$

The second possibility is that the firm undercuts the add-on price but decides against unshrouding the add-on, leaving the fraction of myopic consumers at α. In this case, the deviant firm will only attract sophisticated consumers because myopes do not take notice of the change in the add-on price. The deviation profit in this case is equal to:

$$\pi_i^{dev} = \left[\frac{\alpha}{n} + (1-\alpha)(1 - F(\hat{p}_i^{dev}))\right] \hat{p}_i^{dev}. \tag{1.5}$$

[11]This result follows from the assumption that the inner solution was optimal in the collusive play ($\frac{\overline{v}}{4(1-\alpha)} \geq \alpha\bar{p}$).

[12]Note that it suffices to consider the case $\hat{p}_i^{dev} \leq \overline{v}$: If a deviating firm decides to unshroud the add-on, all consumers are sophisticates, who only consider to purchase the add-on if $\hat{p}_i^{dev} \leq \overline{v}$. Hence, deviating with the corner solution cannot be optimal.

If a firm deviates and decides not to educate consumers, there again could be an 'inner' solution and a 'corner' solution. The 'inner' solution is the case for which $1 - F(\hat{p}_i^{dev}) > 0$. Maximization of (1.5) then yields:

$$\hat{p}_i^{dev} = \frac{\overline{v}}{2(1-\alpha)}\left(1 - \alpha\frac{n-1}{n}\right),$$

yielding a deviation profit of

$$\pi_i^{dev} = \frac{\overline{v}}{4(1-\alpha)}\left(1 - \alpha\frac{n-1}{n}\right)^2. \tag{1.6}$$

The 'corner' solution of the deviation is the case for which $1 - F(\hat{p}_i^{dev}) = 0$. In this case it is optimal to set $\hat{p}_i^{dev} = \bar{p}$. The deviating firm would attract no additional consumers. It would only attract myopic consumers such that the total demand for the firm would be $\frac{1}{n}\alpha$. This deviation might yield a higher profit than (1.6) if the upper bound on the add-on price is extremely high. The add-on price $\hat{p}_i^{dev} = \bar{p}$ would then yield a deviation profit of

$$\pi_i^{dev} = \frac{\alpha\bar{p}}{n}. \tag{1.7}$$

Closer inspection of (1.7) yields that $\hat{p}_i^{dev} = \bar{p}$ can never be a profitable deviation. First note that playing $\hat{p}_i^{dev} = \bar{p}$ and deciding not to educate consumers only corresponds to an actual deviation from collusive behavior if the colluding firms played the inner solution with profits of $\pi^M_{inner} = \frac{\overline{v}}{4(1-\alpha)}$. Then it must be the case that these profits are higher than the profits for the corner solution $\pi^M_{corner} = \alpha\bar{p}$. It follows directly that the deviation profit (1.7) must be lower than the shared collusive profits.

To sum up, if a deviation is profitable, a deviating firm that unshrouds the add-on earns a profit of $\pi_1^{dev} \equiv \frac{\overline{v}}{4}$. The maximum profit that a deviating firm can obtain if it does not unshroud is $\pi_2^{dev} \equiv \frac{\overline{v}}{4(1-\alpha)}\left(1 - \alpha\frac{n-1}{n}\right)^2$. As we formally show in the Appendix it holds that $\pi_1^{dev} > \pi_2^{dev}$. Hence, the following result applies:

Proposition 1.2. *If a profitable deviation exists, a firm that deviates from collusive play will unshroud the add-on.*

As we have mentioned above, there exist two opposing effects of unshrouding. On the one hand, unshrouding increases the number of sophisticated consumers who notice the deviation. On the other hand, unshrouding crowds out formerly myopic consumers. At first glance it was not clear which of these effects is generally stronger, but now we can argue that the positive effect dominates the negative one.

In the following we want to determine the effect that the existence of myopic consumers has on the stability of collusion. To do this we determine the critical discount factors. If firms play an inner solution in the collusive phase, the critical discount factor is given by:

$$\delta^{inner} = \frac{n - \frac{1}{1-\alpha}}{n},$$

which is falling in the share of myopic consumers α. If collusion was characterized by a corner solution, the critical discount factor takes the form:

$$\delta^{corner} = \frac{n - \frac{4\alpha\bar{p}}{\bar{v}}}{n},$$

which is also falling in α. Now we want to show that the critical discount factor is globally falling in α. Since δ is falling piecewise, it suffices to show that δ has no 'jump' when the optimal monopoly strategy changes from the inner solution to the corner solution. At the point of indifference between inner solution and corner solution, it holds that $\frac{\bar{v}}{4(1-\alpha)} = \alpha\bar{p}$. Then it immediately follows that for this parameter constellation, it holds that $\delta^{inner} = \delta^{corner}$. This suffices to ascertain that the critical discount factor δ is continuous in α. Since δ is also falling locally in α, we can conclude:

Proposition 1.3. *The critical discount factor δ is globally falling in the fraction of myopic consumers α.*

The intuition behind this finding lies in the fact that a deviation from collusion is less rewarding with the existence of myopic consumers. First, recall that monopoly profits and, with it, individual collusion profits are increasing in the fraction of myopic consumers. This is because myopic consumers underestimate the price of the product bundle and therefore are more likely to buy it. Hence, for given prices, the total demand is increasing in α. Second, we have shown that a deviating firm optimally unshrouds the add-on. Since then all consumers are sophisticated, total demand is independent of α and for given prices lower than in the collusive phase. Hence, the higher the initial share of myopic consumers, the less attractive a deviation gets when compared to the collusive play, and therefore the more stable is collusion.

Up to now, we assumed that firms always coordinate on monopoly profits and then showed that collusion is more stable the higher the share of myopic consumers is. Beside the number of firms n, it is the collusion-to-deviation profit ratio that determines the critical discount factor and hence, the stability of collusion. When the market contains only sophisticated consumers, it is obviously optimal to collude with the monopoly prices since a deviating firm can always earn at least the aggregate collusive profits and hence, collusion can not be stabilized by coordinating on other than monopoly profits. This does not hold if some consumers are initially myopic. A deviating firm is then not able to attract the whole demand by undercutting marginally since either the existing myopic consumers shop randomly or unshrouding leads to a crowding-out of formerly myopic consumers. Hence, it is not obvious at first glance that the above described ratio is maximized by coordinating on monopoly profits. Collusion could then possibly be stabilized by coordinating on other than monopoly profits such that a deviation gets less attractive. Yet if the market contains myopic consumers, it still holds that collusion is most stable if firms coordinate on monopoly profits.

Proposition 1.4. *Collusion cannot be stabilized by coordinating on other than monopoly profits.*

As we have shown above, the critical discount factor δ is falling in α. The central implication of this result is that the higher the share of myopic consumers α is, the easier it is to sustain collusion. This implies that firms may have an incentive to raise the fraction of myopic consumers. If, initially, there does not exist a collusive equilibrium, active shrouding can potentially decrease the critical discount factor sufficiently, such that collusion becomes sustainable. In addition, shrouding results in higher collusive profits and hence is beneficial for firms even if collusion is already sustainable without additional obfuscation.

We can now conclude that, in a market with a positive fraction of myopic consumers, the critical discount factor is always strictly lower than it would be if the whole consumer population was sophisticated. This follows directly from the result that the critical discount factor is globally falling in α.

Corollary 1.1. *The existence of myopic consumers facilitates collusion.*

Another result that directly follows from closer inspection of the critical discount factors is the following:

Corollary 1.2. *The critical discount factor rises with the number of firms.*

The result that collusion is less stable if the number of firms rises is not particularly new. However, our results indicate that, if the number of firms has an impact on the stability of collusion, the number of firms also has an impact on the whether or not firms shroud the add-on. This is because a breakdown of the collusive shrouding equilibrium leads to an unshrouding of the market. One study that supports this result is Miravete (2007), who presents a study about the U.S. cellular telephone industry. He shows that the entry of new firms to the market tends to 'lift the fog' and leads to more transparent pricing schemes. This finding is in line with our results that an increase in the number of firms may destabilize collusion, which in turn leads to unshrouded and transparent markets.

1.3.3. Welfare Analysis and Regulatory Intervention

We will now further analyze the welfare implications of regulatory intervention. We are interested in the effects that the firm behavior, in particular shrouding the add-on, has on welfare. We argued before that the regulator may want to prevent collusion and shrouding, but the exact effects on welfare have not been thoroughly derived yet. In line with many previous authors like O'Donoghue and Rabin (1999), we deem the true consumer utility to be the relevant measure of consumer welfare. This true consumer utility stands in contrast to the anticipated consumer utility that we interpret as determining the choice of consumers, but as having only a distorted connection to the real utility of consumers.[13]

The total welfare is simply the sum of the valuations of all consumers that buy the product bundle. The price that consumers pay for the product bundle has no impact on welfare because it simply is a redistribution from consumer surplus to industry profit. If all consumers buy the product bundle, for example if firms play the unshrouded competitive equilibrium, welfare would be $E[v] = \frac{\overline{v}}{2}$. If, however, firms play the shrouded collusive equilibrium, not all consumers buy the product bundle, which results in a lower total welfare. Hence, the following Lemma applies:

Lemma 1.2. *The total welfare is higher in the unshrouded competitive equilibrium than in the shrouded collusive equilibrium.*

The result that shrouding is detrimental for welfare is in line with the findings of Gabaix and Laibson (2006) and Kosfeld and Schüwer (2011). Note, however, that their results stem from the assumption that there exists a substitution that can replace the add-on. In particular, these papers assume that the substitution is lost in terms of welfare. However, in many cases it seems plausible that some part of its price is not completely lost.[14] In this case, total welfare would be independent of the shrouding decisions of firms.

Now we want to determine the effects of firm behavior on the consumer surplus. Obviously, the case in which firms charge competitive prices is better for consumers than the case in which firms charge monopoly prices and shroud the add-on. Also note that firm profits are increasing in the fraction of myopic consumers, which is to the detriment of the consumer surplus. In addition, the higher the fraction of myopic consumers, the more consumers buy the product although it yields negative utility to them. Hence, the following Lemma applies:

Lemma 1.3. *The consumer surplus is higher in the unshrouded competitive equilibrium than in the shrouded collusive equilibrium. The consumer surplus is falling in the fraction of myopic consumers α.*

The derived results have wideranging implications for regulatory policies. The most obvious regulation would be to force firms to offer their products at zero prices. Needless to say, this might not be enforceable in real world markets. However, there are other kinds of regulations that can also have positive effects on welfare. For example, our results give new insights into the usefulness of regulations with regard to consumer education and market transparency.

[13]This implies that myopic consumers do not act according to their own best interest. In other words, their myopia is not a sign of different taste, but a sign of a particular malfunction of their decision behavior.

[14]Consider, for example, the case of hotel telephones. If a consumer brings her cell phone with her to save the costs of the hotel line, the calling costs for the cell phone are not lost in terms of welfare, but are part of third party firm profits.

The traditional reason for such regulations was that these should enable consumers to make wiser purchase decisions, which in turn was supposed to increase consumer welfare. Our paper presents another reason for such regulations that, in terms of the model, are intended to reduce the fraction of myopic consumers. If the regulator can unshroud the add-on and thereby force firms to play the unshrouded competitive equilibrium,[15] both consumer surplus and total welfare can be increased. Note here that a regulatory intervention is also increasing the consumer surplus if the intervention cannot decrease the fraction of myopic consumers to zero. First, a regulatory intervention that decreases α also increases the critical discount factor δ^*, which might make collusion infeasible. Second, even if collusion still is stable, a regulatory intervention nevertheless increases consumer welfare since the consumer surplus decreases in the fraction of myopic consumers.

Another innovation of our model lies in its predictions on the detection of cartels. Traditional competition policy had to watch out for active arrangements or coordination between firms in order to detect collusive behavior. Our model predicts that agreements on prices may not be the only sign of collusion. We have argued that the level of obfuscation may be artificially increased by firms in order to stabilize collusion. Hence, the regulator can use the degree of obfuscation in a market as a proxy for the degree of cartelization. This relation seems a useful extension to traditional antitrust monitoring since obfuscation is usually far easier to detect than active coordination between firms.

Apart from active consumer education, the regulator can also intervene by reducing barriers to entry for new firms or by employing other measures that increase the number of firms that participate in the market. This may inhibit collusion because the critical discount factors for both the inner and corner solution are increasing in the number of firms.

1.4. Conclusion

We have proposed a model of limited attention in which competitive firms can either shroud or unshroud the add-on market. We have shown that two kinds of equilibria exist. In one equilibrium firms collude on monopoly pricing and shroud the add-on. In the other equilibrium firms set prices competitively and unshroud the add-on. The equilibrium in which firms shroud the add-on is only stable if the discount factor of the firms is above a critical discount factor. It turns out that this critical discount factor is decreasing in the fraction of myopic consumers. Hence, firms might try to increase the degree of obfuscation and thereby increase the fraction of myopic consumers, which in turn will tend to stabilize collusion. Another incentive to increase

[15] For example, if the regulator can decrease the fraction of myopic consumers α to zero, the critical discount factor increases to $\delta^* = \frac{n-1}{n}$. Hence, if the common discount factor of firms is not sufficiently high, unshrouding the add-on is likely to destabilize collusion.

the level of obfuscation is that the profit that firms earn when they collude is increasing in the fraction of myopic consumers.

These results suggest several implications for welfare analysis and regulatory intervention. We find that welfare is maximized if firms do not collude and do not shroud the add-on. Hence, the regulator might employ measurements to increase consumer sophistication. If these measurements are sufficiently efficient, the collusive equilibrium breaks down and only the natural competitive equilibrium remains.

Our results also suggest new insights into competition policy. We have shown that shrouding might be used by firms as a tool to stabilize collusion. Hence, the degree of obfuscation in a market might be a proxy for the degree of collusion and hidden industry agreements. Markets with high obfuscation then are candidates for further investigation by the antitrust agencies.

2. Gender Differences in Competition and Sabotage

2.1. Introduction

Although they make up nearly half of the workforce, it is a well-known fact that females are underrepresented in upper hierarchy levels of companies worldwide. In January 2012, around three percent of the largest publicly listed companies in the European Union had a female president or chairperson, and the share of females on corporate boards was 13.7 percent. A similar pattern can be observed in the United States: in January 2013 only 21 CEOs of the Fortune500 companies were female, resulting in a share of 4.2 percent.[1] Researchers have offered several explanations for this fact such as labor market discrimination, differences in education, preferences, or biological factors as well as the reluctance of females to enter competitions (e.g., promotion tournaments). Many studies have shown that male participants react more strongly to competitive incentives (e.g., Gneezy, Niederle, and Rustichini (2003)) while females have a tendency to abstain from competition and prefer wage schemes with absolute instead of relative compensation (see Croson and Gneezy (2009) for an overview). Furthermore, there exists evidence that good performance in a tournament and winning per se has a stronger impact on self-esteem of males than on self-esteem of females (see, e.g., Crocker, Luhtanen, Cooper, and Bouvrette (2003) or Wieland and Sarin (2012)). One aspect that has not been discussed in the literature investigating gender differences in tournaments is the question if males and females differ in their willingness to destroy output and sabotage their opponents to ensure winning the competition, even though sabotage is not rare in organizations and tournaments are especially prone to such behavior. In promotion tournaments, sabotage might lead to selecting the less able candidate for a promotion. If, for instance, males have a higher willingness to sabotage because they react more strongly to competitive incentives or enjoy winning per se, as indicated by pre-

[1] See http://ec.europa.eu/justice/gender-equality/files/women-on-boards_en.pdf for the European and http://www.catalyst.org/knowledge/women-ceos-fortune-1000 for the US data.

vious findings mentioned above, this might help to explain why females are underrepresented at leading positions or refrain from entering a competition in the first place.

We fill this gap in the existing literature and study the actions of males and females in a tournament with sabotage opportunities. For this purpose, we conducted a real-effort experiment where two players participated in a rank order tournament and had the opportunity to sabotage each other by destroying a certain amount of work of their opponent. To come closer to real world situations, we introduced a principal to our setting who was paid based on the output of the contestants. Hence, sabotaging not only affected the opponent but also reduced the payment of the principal. We conducted four different treatments: the baseline (as described above), belief, cheating, and gender treatment. In the belief treatment, we elicited beliefs about the performance in the real effort task as well as the chosen sabotage of the respective opponent to analyze if those beliefs differ between males and females. The cheating treatment allows us to check whether social preferences with respect to the principal affected the contestants' behavior. In the gender treatment, we revealed the gender of the opponent before the tournament to study if the contestant's behavior depends on the gender of the respective opponent. Our main findings can be summarized as follows: We find that (i) the males and females on average performed equally well in the real effort task (except for the gender treatment) but (ii) males chose significantly higher levels of sabotage than females. Males were therefore much more likely to win in tournaments with mixed gender participants. Despite this difference, males and females received similar payments because sabotage was costly. The gender gap is not only present in actual sabotage choices, but in the stated beliefs about the opponent's actions as well. Males not only sabotaged their opponent more severely, they also expected their opponents to inflict more sabotage on them. If we revealed the gender of the opponent, we also observe a gender gap in performance. Males performed significantly better in the real effort task than females. In the sabotage dimension, both females and males believed to be sabotaged more severely from males, but we do not find any differences in sabotaging behavior with respect to the revealed gender of the opponent. Our main finding, the gender gap in sabotage, was persistent over all treatments and cannot be explained for instance by differences in risk attitudes, human values, or social preferences with respect to the principal.

2.2. Related Literature

Our paper is related to the literature on sabotage in tournaments as well as on gender differences in competition. In his seminal paper, Lazear (1989) shows that the optimal wage spread is lower when participants are able to sabotage each other. Hence, the tournament designer optimally uses a more equitable prize structure in order to lower the incentives to sabotage the opponent.[2] Because company data on sabotage is generally not available for research, empirical studies use sports data (see, e.g., Garicano and Palacios-Huerta (2006), del Corral, Prieto-Rodriguez, and Simmons (2010), Balafoutas, Lindner, and Sutter (2012), or Deutscher, Frick, Gürtler, and Prinz (2013)) or experimental data to investigate the impact of sabotage on tournaments. Harbring and Irlenbusch have contributed several papers dealing with different prize spreads, a varying number of participants, and different numbers of tournament prizes (see, e.g., Harbring and Irlenbusch (2008)) as well as communication in tournaments with the possibility to sabotage in lab experiments (Harbring and Irlenbusch (2011)).[3] While most papers use a chosen effort setting, we are aware of only a few papers that implement a real-effort tournament with sabotage which are closer to our study. Vandegrift and Yavas (2010) use a forecasting task and give the contestants the option to raise the forecasting error of their opponent. Players do not know the performance of their opponent when selecting the costly sabotage. The cost function of sabotage is linear meaning that the players have to pay a constant fee for each additional unit of sabotage. Players exert more sabotage if the prize spread is higher or the players are rematched after each period. In the study of Carpenter, Matthews, and Schirm (2010), the task was to prepare letters and envelopes. The authors conducted different treatments with piece rate and tournament incentives, as well as with and without sabotage. They find that output declines in the tournaments with sabotage compared to treatments with piece rate. The reduction is due to false reporting of the quality rather than the quantity of the output. Hence, the players preferred the more subtle form of sabotage and refrained from "undercounting" the total output of an opponent. In contrast to our paper, the players in the setting of Carpenter, Matthews, and Schirm (2010) selected their amount of sabotage after the production period when they already knew their own performance. Hence, Carpenter et al. studied a sequential tournament. Recently, Charness et al. (forthcoming) matched players into groups of three and let them work on a decoding task in a flat wage environment. In the absence of monetary incentives, ranking feedback leads individuals to invest in costly sabotage in order to improve their relative position in the group. In contrast to our experiment, there are no monetary incentives to sabotage the opponent and players decide about sabotage after being informed about their relative performance. Note that neither Vandegrift and Yavas (2010) nor Carpenter, Matthews, and Schirm (2010) explicitly report results regarding controls for gender while Charness et al. (forthcomining) mention in a footnote that they do not find gender differences in their setting. Our work differs from these

[2] For further theoretical work on sabotage in tournaments, see, among others Chen (2003), Kräkel (2005), Münster (2007), or Gürtler (2008).

[3] For further experimental evidence on tournaments with sabotage, see, e.g., Harbring, Irlenbusch, Kräkel, and Selten (2007), Falk, Fehr, and Huffman (2008), or Gürtler, Münster, and Nieken (2013) or the recent surveys of Dechenaux, Kovenock, and Sheremeta (2012) , Chowdhury and Gürtler (2013) and Amegashie (2013).

papers because we implemented a principal in our setting which induced a negative externality of sabotage on an uninvolved player, and we had convex costs of sabotage which were identical for all agents. Furthermore, our study concentrates on the impact of gender on sabotage decisions, which is why we conducted different treatments with and without revelation of the opponents' gender as well as positive and negative externalities.

Our work is as well related to the growing literature on gender differences in competition. One strand of this literature analyzes the entry decision and studies the question whether females prefer different incentive schemes than males when they are able to choose their compensation scheme. Niederle and Vesterlund (2007) show that females, compared to equally able males, mostly refrain from competition and instead select a piece-rate scheme, whereas the majority of male participants enters the tournament. This self-selection effect can be considered to be very robust since this finding was replicated by several studies using a similar setup (see, e.g., Wozniak, Harbaugh, and Mayr (2010), Healy and Pate (2011)) as well as by other studies using different designs (see, e.g., Vandegrift and Brown (2005), Gneezy, Leonard, and List (2009), Dohmen and Falk (2011), Price (2012), Buser (2012) and Garratt, Weinberger, and Johnson (2013)).[4]

Another strand of this literature studies the question whether females and males react differently to competitive payment schemes such as rank order tournaments. Gneezy, Niederle, and Rustichini (2003) conducted an experiment in which participants had to carry out a task (solving mazes), being paid according to different compensation schemes. They find no significant performance difference between females and males under a piece-rate scheme, whereas under a competitive compensation scheme, males solved significantly more mazes than females. This finding was replicated by Gneezy and Rustichini (2004) in a field study on the competitiveness of ten year old children. Similar effects can be found in Masclet, Peterle, and Larribeau (2012) who report that males exert higher levels of effort in a competitive environment than females with similar ability. Although evidence points toward males reacting more strongly to competitive incentives than females, this finding also seems to depend on the kind of task. Günther, Ekinci, Schwieren, and Strobel (2010) report no performance differences given a gender-neutral task while they find females to perform better than males in a "female" task. The authors explain their finding with the so-called stereotype threat. Further evidence comes from Iriberri and Rey-Biel (2012) who also find that females only perform worse than males if they believe the task is one where males have an advantage.[5] The gender composition of the group seems to be

[4]The decision to enter a tournament could also be driven by overconfidence. Several studies report that males are generally found to be more overconfident than females. See, e.g., Lundeberg, Fox, and Punćcohaŕ (1994), Beyer and Bowden (1997), Barber and Odean (2001), Bengtsson, Persson, and Willenhag (2005), or Reuben, Rey-Biel, Sapienza, and Zingales (2012).

[5]See Barankay (2012) for gender differences regarding the reaction to ranking information which is not tied to compensation.

important as well: in Gneezy, Niederle, and Rustichini (2003), the gender gap in performance is significantly higher in mixed-sex than single-sex tournaments. Moreover, Datta Gupta, Poulsen, and Villeval (2013) report that males tend to compete less against males than against females, which might also be affected by the afore-mentioned self-esteem effects.[6] Further evidence that males tend to perform slightly better if they compete against a female opponent comes from Antonovics, Arcidiacono, and Walsh (2009), Price (2008), Iriberri and Rey-Biel (2012), or Cotton, McIntyre, and Price (2013),.

While all these papers investigate productive behavior in competitive situations, we are aware of only one paper which also studies gender effects with respect to unethical behavior in tournaments. Schwieren and Weichselbaumer (2010) allow contestants to cheat in order to win a tournament.[7] The players had to solve mazes and in order to cheat could either use an "auto solve option" or a "path verifying option" (which showed them if the chosen path was wrong). With a spy software the authors were able to detect the cheating while the players were not aware that their actions would be observed. Cheating was not costly for the players and the authors find that cheating depends on the performance of a contestant. Lower performing players cheated more than better players. Even though females on average performed worse than males one cannot conclude that females in general are also more likely to cheat. Hence, the observed gender gap in cheating is driven by differences in performance. In contrast to Schwieren and Weichselbaumer (2010), we studied sabotage decisions in tournaments where sabotaging was costly. In our setting, exerting sabotage not only harmed the opponent but also resulted in a lower payment for the principal.[8]

2.3. Experimental Design

We implemented a simultaneous two-player real-effort tournament with the option to sabotage the opponent. The experiment encompassed four different treatments: baseline, belief, cheating, and gender treatment (see Table A.1 in the appendix for an overview). We start our description with the baseline treatment, consisting of 8 working periods which lasted 5 minutes each.[9] One of those periods was randomly selected for payment. In each session, we had 21 participants who were divided into three units. Each unit contained one principal and six agents.

[6]For results of team tournaments see, e.g., Apesteguia, Azmat, and Iriberri (2012), or Dargnies (2012).

[7]For further research on cheating in tournaments see, e.g., List, Bailey, Euzent, and Martin (2001), Preston and Szymanski (2003), Enders and Hoover (2004), Shleifer (2004), and Charness et al. (forthcoming).

[8]For an extensive overview of gender and competition, see Niederle and Vesterlund (2011) or Croson and Gneezy (2009) for a more general survey of gender differences in preferences. Furthermore, Dechenaux, Kovenock, and Sheremeta (2012) provide a recent survey of experimental results in contests and tournaments. For details on gender see chapter 8.7.

[9]The instructions of the baseline treatment translated into English can be found in the appendix.

Each agent was matched with a different agent from his unit at the beginning of every period. We implemented this stranger matching to prevent reputation effects and reciprocity among the agents.[10] Before the first period started, the participants had the opportunity to get used to the real-effort task in a trial period that was not payoff relevant. At the beginning of every period, each agent had to decide how much sabotage to inflict on the opponent by selecting an integer $x \in [0, 70]$ which would be deducted from the other contestant's achieved points. Exerting sabotage (destructive effort) was costly, because in reality sabotaging an opponent tends to result in opportunity costs of time, the contestant might need to exert some extra effort to cover up his destructive activity, and there is always the danger of being detected and punished. The cost function was given by $c(x) = x^2/14$. Next, agents had to carry out the real-effort task in the five-minute working period. In contrast to chosen effort, real work "involves effort, fatigue, boredom, excitement, and other affections" (Van Dijk, Sonnemans, and Van Winden (2001), p. 189). Regarding the decision to sabotage and destroy performance, we believe it is important that performance is based on real work rather than on a chosen number. In line with Carpenter, Matthews, and Schirm (2010) and Vandegrift and Yavas (2010), we therefore, preferred a real-effort setting. Similar to Erkal, Gangadharan, and Nikiforakis (2011), the participant's task was to encode words as numbers. They were asked to enter a two-digit number for each letter of a word according to an encryption table that assigned a number to each letter of the alphabet. Agents received one point for each correctly encoded letter and could proceed to encoding the next word only if all letters had been encoded correctly. Note that all participants received the same words in the same order. The points earned by encoding were summed up after each period and we refer to them as performance in the remainder of the paper. An agent's output was given by the achieved points minus the suffered sabotage. At the end of each period, the agent with the higher output received 500 taler while the agent with the lower output received 200 taler. In case of a tie the winner was determined by a random draw of the computer. After each period, agents received information on the achieved points in the encoding task (performance), their own choice of sabotage, and whether they had won the tournament or not. Note that they did not learn their own total output (achieved points minus sabotage of the opponent) or the total output of their opponent at any time during the experiment. Hence, they were not able to deduct the amount of sabotage inflicted on them from the information they received after each period. In every period, agents had to decide how much effort to exert and how much sabotage to inflict on the opponent. In our setting, both choices had to be made without intermediate information about the achieved points and the sabotage level of their opponent. The contestants, therefore, could not condition their actions (effort and/or sabotage) on the behavior of their opponent as it would be possible in a sequential setting (see, e.g., Carpenter, Matthews, and Schirm (2010),

[10] Note that our design is not a perfect stranger matching because agents could play against each other more than once. However, they did not know if and when they would meet again.

and Gürtler, Münster, and Nieken (2013)) In reality, we observe both, simultaneous and sequential tournaments. While situations where opponents destroy work after observing the actions of their contestants have the character of a sequential tournament, sales contests or promotion tournaments where the contestants work in different geographical regions or branches resemble simultaneous tournaments where the contestants have no information about the actions of their opponents. Our design is similar to Harbring and Irlenbusch (2008), Vandegrift and Yavas (2010), Harbring and Irlenbusch (2011), or Gürtler, Münster, and Nieken (2013) where contestants also select effort and sabotage simultaneously but in contrast to this paper in a chosen effort setting. The principal's payment was determined by the average output of the six agents in the corresponding unit. The principal observed the output of each agent in her unit but not the amount of sabotage or the points achieved in the encoding task (performance). The agents were aware that their output determined the payment of the principal and they knew which information she was given. We chose to include a principal in the design to study a situation where sabotage not only reduced the chances of the opponent to win the tournament but also affected a third party which was not directly part of the competition. In reality, sabotage usually harms the firm and we wanted to capture this effect in our design. Furthermore, the real effort task of encoding words as numbers might be rather meaningless if it is carried out without a principal whose payment actually depended on the realized output. After the experiment, all participants filled out a questionnaire containing questions about the experiment. Additionally, we collected socioeconomic information, basic human values and elicited the participants' risk attitudes using a question from the GSOEP. This question elicits the general willingness to take risks on a 11-point scale. We used the Schwartz Human Values Questionnaire as implemented in the European Social Survey with 21 items to measure basic human values. According to Schwartz (2000) these values are recognized across different cultures and societies. These are power, achievement, hedonism, stimulation, self-direction, universalism, benevolence, tradition, conformity, and security. For a detailed discussion about human values and gender differences, please see Schwartz and Rubel (2005) and Adams and Funk (2012).

The belief treatment is identical to the baseline treatment, the only difference being that we elicited agents' beliefs about the achieved points (performance) and the amount of sabotage of their respective opponent. In each period, agents were asked to estimate their opponent's achieved points (performance) and sabotage decision before choosing their own amount of sabotage. The elicitation of beliefs was incentivized, agents received $\max\{15 - Z, 0\}$ taler for every stated belief being Z points away from the correct value.[11] In the gender treatment, again

[11]Eliciting beliefs might potentially lead to hedging behavior. Blanco, Engelmann, Koch, and Normann (2010) report that hedging of beliefs in experiments is not "a major problem unless hedging opportunities are very prominent." In our setting, hedging opportunities are not very prominent because a higher output led to a payment of 500 taler while a correct belief led to 15 taler. Furthermore, for instance Kräkel and Nieken (2012) report no differences in the behavior of participants in a tournament setting whether the beliefs have been

everything else being equal to the belief treatment, the gender of the opponent was revealed prior to the elicitation of beliefs. At the beginning of each session, all agents had to state their gender. After the trial period but before each period, each agent was informed about the gender of his or her opponent on a separate screen. After the agent clicked on the OK button, the experiment proceeded. During the whole period the gender of the respective opponent was shown on the top of each screen. The gender of the principal was not revealed.

In the cheating treatment, we switched from a sabotage opportunity to a cheating opportunity. In contrast to the baseline treatment, the chosen amount of points x was now added to the agent's own achieved points in the encoding task instead of being deducted from the opponent's achieved points.

The experiment was conducted at the BonnEconLab. We used the online recruitment ORSEE (Greiner (2004)) and programmed the experimental software in z-tree (Fischbacher (2007)). Each treatment encompassed 4 sessions which lasted 1.5 hours. 336 students enrolled at the University of Bonn participated in the experiment with 192 females and 144 males.[12] The average earnings were 17.49 euro (including an endowment of 6 euro for agents and 4 euro for principals) and the exchange rate was 25 taler for one euro (approx. USD 1.25 at the time of the experiment).

2.4. Results and Discussion

In our analysis, we focus on the behavior of the participants in the role of agents because the principals had no decision power and only received the produced output (achieved points minus sabotage of the respective opponent) in our setting. We start the analysis by checking the achieved points in the encoding task (performance) in the baseline treatment. On average, females achieved 111.43 points while males achieved 114.87 points. We do not find a significant difference in performance between male and female agents ($p = 0.4183$ at subject level; $p = 0.2482$ at session level).[13] Hence, in terms of performance, our results are in line with the

elicited incentivized or not.

[12]Note that we have to drop one subject in the gender treatment because the subject selected female as gender in the beginning and male as gender at the end of the experiment. Because the opponents were under the impression of competing against a female and did not receive information about the performance of the subject, we kept those observations. The results are qualitatively the same if we drop all subjects that were in the same matching group (five other agents).

[13]For all non-parametric comparisons, we report two-sided Mann-Whitney U tests with data pooled over all periods for each agent. In addition, we also report the results for the data pooled at session level to take into account that the observations of the agents in one session are not stricly independent of each other. Note that the subjects were not informed about performance or sabotage levels of their opponents and were matched in groups of six agents and one principal resulting in three matching groups for each session. Hence, we have

findings of Niederle and Vesterlund (2007), Günther, Ekinci, Schwieren, and Strobel (2010), Wozniak, Harbaugh, and Mayr (2010), and Healy and Pate (2011). While we observe similar performance levels of males and females in the encoding task, we find an improvement of performance over the course of the experiment for both genders (see Figure 2.1).

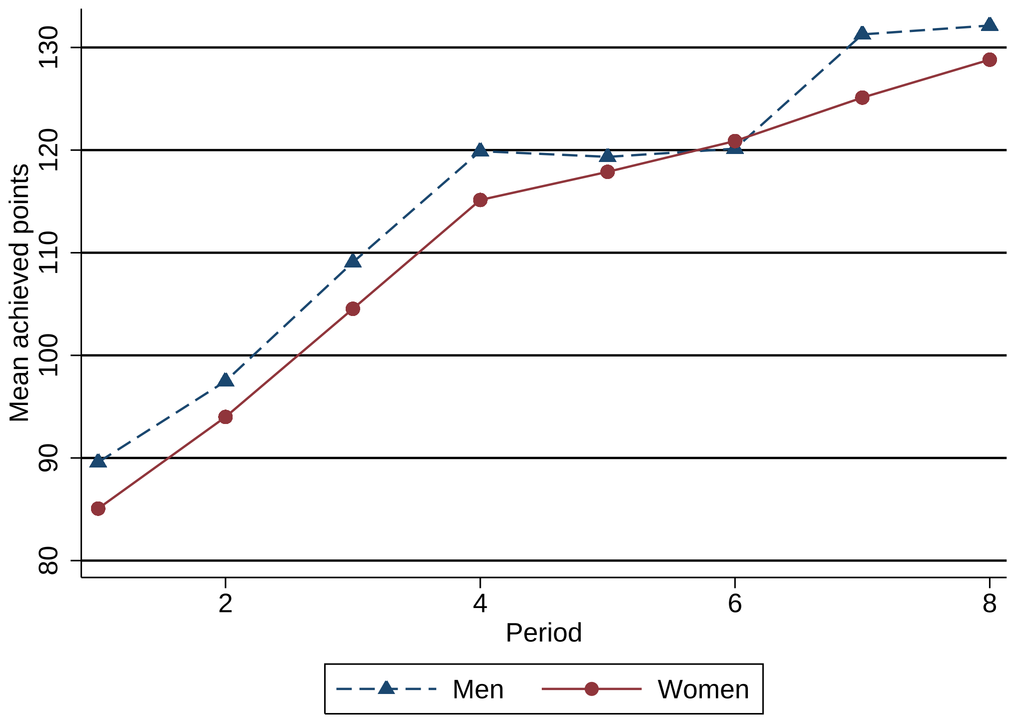

Figure 2.1.: Mean of achieved points (performance) in each period for males and females in the baseline treatment

This is supported by the regressions with the achieved points (performance) as the dependent variable (see columns (2) to (4) of Table A.2 in the appendix).[14] We include a dummy variable for females and control for time trends and risk attitude (lower values indicating a higher risk aversion). Gender has no significant impact on the achieved points, but the variable period controlling for a time trend has a significantly positive impact indicating learning effects. We also included an interaction effect between the female dummy and period to check if males and females differ regarding performance over time but the interaction is not significant.

Next, we examine the sabotage decisions of the agents in the baseline treatment. On average males selected 26.63 points of sabotage while females selected 12.99 points of sabotage. The

three groups within each session that did not interact with each other.

[14]We apply Random Effects GLS regressions to take the panel structure of the data into account and calculate robust standard errors clustered on session level. Additionally, we ran those regressions with robust standard errors clustered on matching group level and subject level. The results remain qualitively the same and can be obtained upon request.

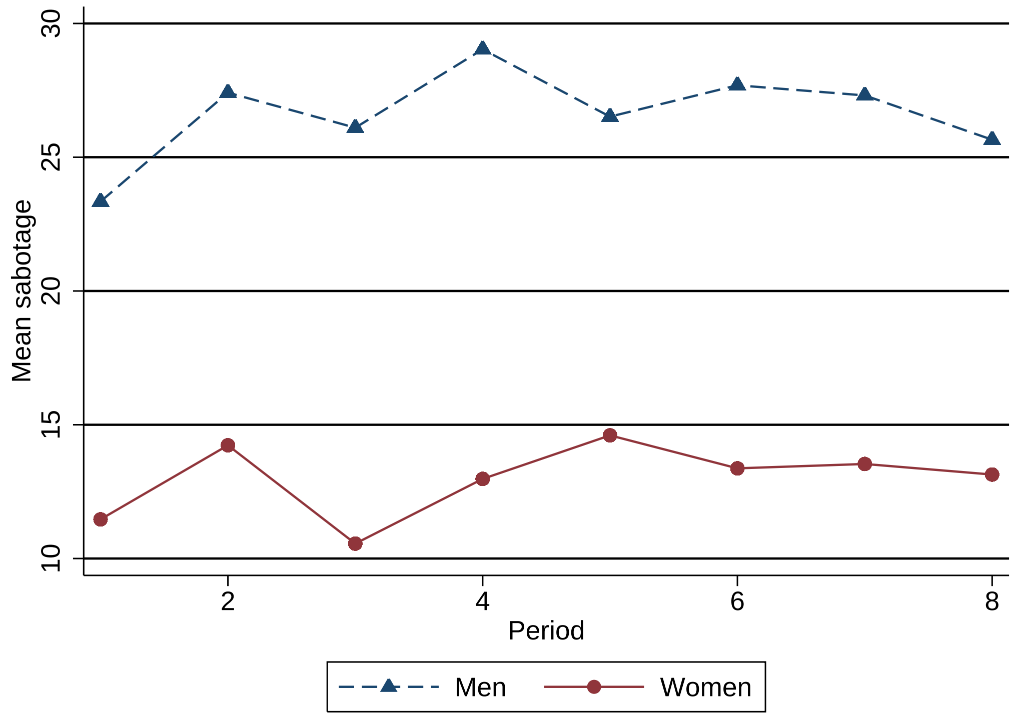

Figure 2.2.: Mean of sabotage in each period for males and females in the baseline treatment

observation that females selected significantly less sabotage than males is supported by Mann-Whitney U tests using pooled data over all periods ($p = 0.000$ at subject level; $p = 0.0209$ at session level). A closer look reveals that we have 120 observations of females who selected zero sabotage, while this is only true for 43 observations of males. Hence, females more often preferred not to sabotage at all. But even if we only compare those observations where a positive amount of sabotage was chosen, males were more prone to larger amounts of sabotage than females (average sabotage for males 32.69 points and for females 19.94 points). We report the results of Random Effects GLS regressions with robust standard errors clustered on session level and the amount of sabotage as the dependent variable, see columns (1) to (3) of Table 2.1.[15] Because the amount of sabotage could range between zero and 70, we also applied Random Effects Tobit regressions (see columns (4) to (6) of Table 2.1). The dummy for females is negative and highly significant in all specifications. As can be seen in Figure 2.2, there is a slight rise of sabotage during the first periods of the experiment. We control for a time trend by including the variable period in the regressions in Table 2.1 but it shows no significant

[15]We also ran those regressions with robust standard errors clustered on matching group as well as subject level. The results are robust and available upon request. Additionally we ran a two part model (see, e.g. Manning, Duan, and Rogers (1987)) showing that we have gender differences both in the decision whether to sabotage or not and the amount of chosen sabotage. The results are available upon request. We thank an anonomous referee for this suggestion.

effect. Also the interaction between female and period is not significant. While we do not find significantly different levels of performance in the encoding task, males selected on average twice as much sabotage than females.

	Random Effects GLS			Random Effects Tobit		
	(1)	(2)	(3)	(4)	(5)	(6)
Dummy female	-13.65^{***}	-13.18^{***}	-13.20^{***}	-18.13^{***}	-17.49^{***}	-16.76^{***}
	(4.014)	(3.845)	(1.965)	(5.177)	(5.237)	(5.724)
Period		0.216	0.213		0.0771	0.173
		(0.219)	(0.453)		(0.255)	(0.394)
Risk attitude		0.530	0.530		0.682	0.683
		(0.400)	(0.400)		(0.988)	(0.988)
Female x period			0.00483			-0.165
			(0.490)			(0.517)
Constant	26.63^{***}	22.85^{***}	22.86^{***}	25.39^{***}	21.43^{***}	20.99^{***}
	(3.557)	(3.073)	(2.629)	(3.989)	(6.681)	(6.819)
Observations	576	576	576	576	576	576
# of left censored				163	163	163
# right censored				21	21	21
Log Likelihood				-1750.0491	-1749.7657	-1749.7149
R^2	0.186	0.143	0.143			

Dependent variable is sabotage. Standard errors (GLS: Robust standard errors clustered on sessions) in parentheses.
$^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

Table 2.1.: Random Effects GLS and Random Effects Tobit Regressions for the baseline treatment with sabotage as the dependent variable

In the following, we discuss several possible explanations for this highly significant gender gap in sabotage. We examine whether these differences can be explained by risk attitudes, basic human values, or aspects of social preferences. Furthermore, we analyze the beliefs about the actions of the respective opponent to see if the gender gap was also present in the beliefs. In a last step we report the data of the gender treatment to see if the behavior changed when the gender of the respective opponent was revealed.

Since previous studies report gender differences in the risk attitude (see, e.g., Holt and Laury (2002), Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner (2011) , and Charness and Gneezy (2012) we checked the impact of the general risk attitude of the agents (measured by the question taken from the GSOEP) on the amount of chosen sabotage but find no significant effect (see columns (2), (3), (5), and (6) of Table 2.1). We do not find any significant differences between males and females in risk attitudes in our sample either ($p = 0.1725$ at subject level). This might be due for instance to the fact, that students are a rather homogeneous group regarding other factors such as intelligence or age which are also known to influence risk attitudes besides gender. Note that we use a question taken from the GSOEP to elicit the general willingness to take risks. As Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner (2011) have shown, it is a good predictor of actual risk-taking behavior. Using a representative sample of 450 German subjects, Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner (2011) find that the survey question predicts risky behavior in an incentivized lottery task and the risky behavior in different situations.[16]

Moreover, differences in basic human values might have influenced the sabotage decisions. We elicited 10 human values using a questionnaire introduced by Schwartz, Melech, Lehmann, Burgess, Harris, and Owens (2001).[17] The ten human values are: Self-direction, stimulation, hedonism, achievement, power, security, conformity, tradition, benevolence, and universalism. According to value theory, values are "desirable, trans-situational goals, varying in importance, that serve as guiding principles in people‘s lives." (Schwartz and Rubel (2005), p. 1010). We expect power and benevolence to be of particular interest in our setting. Winning a competition can have an effect on social esteem and personal welfare. Power focuses on social esteem and is described as "social status and prestige, control or dominance over people and resources" (Schwartz and Rubel (2005), p. 1). We, therefore, presume that agents with a higher level of power are more likely to sabotage in order to increase their chances to win the tournament. Agents with higher levels of benevolence, however, should be less prone to sabotage. Benevolence is described as "preservation and enhancement of the welfare of people with whom one is in frequent personal contact" (Schwartz and Rubel (2005), p. 2). Schwartz and Rubel (2005) have shown that males and females differ in the importance they attach to some human values. While power (as well as stimulation, hedonism, achievement, and self-direction) seems to be of higher importance to males, benevolence (and universalism) is more important to females in the general population. Adams and Funk (2012) confirm these differences in the importance attached to power and benevolence for male and female board members. Hence, based on the findings in previous papers, we expect females to score higher on benevolence but lower on power, resulting in lower levels of sabotage.

Furthermore, it might be worthwhile to also take achievement and security into account. Achievement is described as "personal success through demonstrating competence according to social standards." (Schwartz and Rubel (2005), p. 1). Hence, the effect of achievement on sabotage might be ambiguous. If agents scoring high on achievement are interested in personal success, we would expect a positive impact on sabotage while agents might exert less sabotage

[16]We have also checked interaction effects between gender and risk attitudes in the regressions but again find no significant impact.

[17]Note that we have to drop all observations where agents stated "I do not know" from our sample which leaves us with 536 observations for the regressions instead of 576.

if they feel that sabotaging is not in line with social standards. Security stands for "safety, harmony, and stability of society, of relationships, and of self" (Schwartz and Rubel (2005), p. 2) and therefore agents scoring high on security should be less prone to sabotage. Similar to risk attitudes we find no significant differences regarding power, achievement, and security between males and females in the baseline treatment ($p \geq 0.1368$) and only a weakly significant differences regarding benevolence ($p = 0.0910$) with a higher level for females. As can be seen in column (1) of Table A.3 in the appendix, the basic human values except power have no significant impact on the chosen sabotage. In line with our expectation, higher identification with power leads to more sabotage. However, we have to be careful to draw this conclusion as the effect of power is not robust and vanishes if we apply a random effects tobit regression (see column (3) in Table A.3). Even though in our sample females and males do not differ in the four basic human values with the exception of benevolence, it is worthwhile to study interaction effects between the female dummy and the human values. In both, the random effects GLS and the random effects tobit regression, the interaction between benevolence and female as well as the interaction with security are significantly different from zero. Nevertheless, even after controlling for human values, the female dummy is still significant indicating a persisting non-explained gender gap in sabotage.

Because in our setting sabotage affected the opponent's expected payoff as well as the principal's payoff, differences in social preferences between females and males might help to explain the gender gap in sabotage. First, sabotage reduced the opponent's output and therefore his chances of winning as well as his expected payoff. Hence, social preferences of agents with respect to the opponent might be important. Because a tournament leads to an unequal outcome, the behavior of agents might be influenced by inequity aversion. As Grund and Sliwka (2005a) have shown, in a tournament inequity averse agents select higher amounts of sabotage and exert more effort than self-interested agents if disadvantageous inequity (envy) has a higher impact on utility than advantageous inequity (compassion). While a direct connection to the theoretical results of Grund and Sliwka (2005a) is not straightforward because we have no information about effort cost in our data, we can use their findings as a first guideline. In our experiment, males choose much higher sabotage levels than females. Regarding performance, there is no significant gender gap but a tendency that males tend to perform slightly better than females. In other words, to explain our findings males would have to be far more prone to envy relative to compassion than females. To the best of our knowledge, no empirical results exist showing that this holds true. If males had higher costs of effort than females they might also have tried to compensate for this by sabotaging. However, we have no hints in our data or in the answers given in the questionnaire that there are systematic differences regarding effort costs or ability between males and females. Hence, we do not believe that a gender difference in inequity aversion towards the opponent is the driving factor behind our findings.

Second, sabotage affected the payoff of another, uninvolved player, namely the principal: note that in our setting, similar to many real world situations, sabotage led to a negative externality. Decreasing the output of the opponent not only led to a relative advantage in the tournament but also reduced the payoff of the principal as sabotage destroyed output. If females were less selfish and cared about the principal's payoff to a higher degree than males, they should have chosen lower levels of sabotage. As several experimental results indicate that females are less selfish than males (see, e.g., Eckel and Grossman (1998), Andreoni and Vesterlund (2001), Güth, Schmidt, and Sutter (2007), or Erat and Gneezy (2012)), it might be assumed that gender differences in social preferences with respect to the principal can help to explain the gender gap in sabotage.

We use the results of the cheating treatment to investigate whether females and males differ in their attitude towards the principal. Again, agents had to choose an integer $x \in [0, 70]$ prior to encoding words, but now the chosen amount was added to the agent's own achieved points instead of being deducted from his opponent's achieved points in the encoding task. Note that x affected the agents' payoffs in the same way as in the baseline treatment because it also increased the probability to win the tournament and induced the same cost. The only difference was that cheating was output-enhancing rather than destructive and therefore imposed a positive externality on the principal's payoff. Of course, cheating does not necessarily raise a principal's payoff in reality. Nevertheless, one might consider a positive manipulation of output by an agent as beneficial for the principal, at least in the short run. For instance, doping in sports contests leads to better performance, which most likely attracts more spectators in the short run (as long as the doping decision is private information). This in turn might lead to higher profits for the tournament organizer due to higher revenues from sponsorship and sales of broadcasting rights. Furthermore, imagine a tournament between two sales persons. If one sales person cheats and lies to the potential customers, for instance about the date of delivery or certain features of the product that are not testable (e.g., made of organic materials), and thereby sells more than otherwise, this will at least in the short run be beneficial for the principal. Another example is the manipulation of earnings in a cost center of a big company which makes the company look healthier. As long as the earnings manipulations are not illegal but due to some leeway and "creative" accounting technics this will also be beneficial for the principal. If a systematic difference in the attitude towards the principal between females and males was the driving force behind the observed gender gap in sabotage in the baseline treatment, we should observe a reversed pattern in the cheating treatment with females choosing a higher amount of x than males. In other words, if it was the externality that determined the gender gap in sabotage, a reversal of the externality's sign should affect the choices of x. Another possible line of thought is that agents might perceive cheating to be less nasty than sabotaging because it does not destroy work and therefore cheating has lower moral costs. However, cheating also reduces the chances of the opponent to win, hence sabotage and cheating both "hurt" the opponent to the same extent. Furthermore, Abbink and Herrmann (2011) for instance do not find gender differences regarding the moral costs of being "nasty."

Again, we find no significant difference ($p = 0.8579$ at subject level; $p = 0.7728$ at session level) regarding the achieved points in the encoding task between males (on average 107.03 points) and females (on average 111.83 points). Note that the performance is similar to the results in the baseline treatment ($p = 0.4525$ at subject level; $p = 0.4622$ at session level) but in the cheating treatment males achieved on average less points then females while it is the other way around in the baseline treatment. The average amount of cheating is 25.38 points for male and 14.86 points for female agents, which again is a significant difference ($p = 0.002$ at subject level; $p = 0.0433$ at session level). Similar to the baseline treatment, males selected significantly higher amounts of cheating than females (see also Table A.4 in the appendix). Moreover, we find no significant differences when comparing cheating and sabotage activities in the baseline and cheating treatment, neither for females ($p = 0.6526$ at subject level; $p = 0.7728$ at session level) nor for males ($p = 0.9873$ at subject level; $p = 1.000$ at session level). Hence, the externality on the principal's payoff seems to be irrelevant for the agents' choice of x. In a setting with three players, two competing agents and one principal, social preferences can be quite complex and we cannot rule out that some kind of difference in social preferences between females and males might help to explain the gender gap in sabotage. Nevertheless, the results of the cheating treatment shows that females and males do not differ in their attitude towards the principal, which could have been a possible explanation for the diverging sabotage decisions.[18]

Another driving factor might lie in the expectations about the respective opponent's actions and, based on this, the perception of his relative performance. For instance, if females believed that their opponent's performance was substantially lower than their own, there would have been no need for them to invest in costly sabotage to win the tournament. On the other hand, different expectations about the amount of sabotage by the opponent might have driven males to select higher amounts of sabotage. For instance, agents might select sabotage based on a preemptive retaliation motive and select higher levels to "punish" the opponent. The results of the belief treatment allow us to analyze the agents' estimations regarding the achieved points and the sabotage of the opponent.

First, we compare the achieved points as well as the sabotage decisions with the results of the baseline treatment to see whether the incentivized elicitation of beliefs before each period affected the competition. On average, females achieved 105.06 points and males 110.06 points

[18]Note that the principals on average earned less than the agents in the baseline and the cheating treatment and the payoff difference was lower in the baseline than in the cheating treatment. Hence, agents caring for equal payoff should have chosen less cheating than sabotage. However, the agents were not informed about the payoff of the principal.

in the encoding task. There is no significant difference in performance between male and female agents ($p = 0.4402$ at subject level; $p = 0.3865$ at session level), and the results are similar to those in the baseline treatment. Concerning the sabotage decisions, females, on average, chose 10.60 points of sabotage, while males selected 25.56. As in the baseline treatment, males chose significantly higher levels of sabotage than females ($p = 0.001$ at subject level; $p = 0.0202$ at session level). Because the results are perfectly in line with those of the baseline treatment, we conclude that the elicitation of beliefs prior to the tournament did not change behavior.

Next we inspect the belief about the opponent's achieved points in the encoding task. We do not find a significant gender difference ($p = 0.335$ at subject level; $p = 0.2482$ at session level). Females expected their opponents to achieve, on average, 117.84 points and males expected their opponents to achieve 102.31. The relatively high belief of females is driven by one participant who stated an average belief of 518.75. The medians for males and females are identical with 110 points.

Regarding sabotage, we do find a significant gender gap concerning the beliefs about the opponent's decision: females expected, on average, 15.36 points of sabotage to be inflicted on them, while males believed that the opponents selected 26.13 points of sabotage. The difference between the beliefs is significant ($p = 0.006$ at subject level; $p = 0.0209$ at session level). While we find no gender gap either in the performance in the encoding task or in the beliefs regarding the performance, we do find a gender gap both in the amount of sabotage and in the beliefs about sabotage.

Hence, given no significant differences in the performance dimension, differing beliefs in the sabotage dimension might help to explain the diverging sabotage decisions: males might have invested higher amounts of sabotage in order to compensate for the higher expected amount of sabotage from their opponent or for a preemptive retaliation motive. To test whether beliefs have an impact on the amount of sabotage, we included beliefs about performance and sabotage in the regressions reported in Table A.5 in the appendix. While the belief about performance has no significant effect, a higher belief about sabotage leads to a higher amount of selected sabotage. This finding supports the argument that agents might select sabotage because of a preemptive retaliation motive. However, one has to be careful to draw such a conclusion because the belief might be biased by a preference for consistency. Agents who selected a high amount of sabotage themselves might also have stated that others inflict larger amounts of sabotage on their opponents in order to justify their own choice. Moreover, when deciding on how much to sabotage the opponent, an agent should not only take into account the belief about sabotage, but also include the beliefs about (the opponent's and own) performance. In a next step, we calculate the amount of sabotage which was necessary for each agent to win the tournament according to the stated beliefs. For some agent i to win, his output needed to (weakly) exceed the output of his opponent j. Recall that the output of agent i was composed of his own performance (denoted by e_i) and his opponent's sabotage decision x_j. Hence, agent i won the tournament if

$$e_i - x_j \geq e_j - x_i$$

holds. When deciding how much sabotage to inflict on his opponent, an agent should therefore take into account the belief about his own performance $b_i(e_i)$ as well as the beliefs about the opponent's performance $b_i(e_j)$ and sabotage decision $b_i(x_j)$. Hence, according to his beliefs, and in order to win, agent i's sabotage decision needed to satisfy

$$x_i \geq b_i(e_j) - b_i(e_i) + b_i(x_j) \equiv \underline{x_i}.$$

Because agents were informed about their own performance after each period, we use their performance of the previous period as a proxy for the expected own performance.[19] Along with the elicited beliefs $b_i(e_j)$ and $b_i(x_j)$, we can calculate the amount $\underline{x_i}$, necessary to at least ensure a tie in the tournament. Because we found no significant gender difference in actual performance or in the beliefs regarding the opponent's performance but a significant gender gap in the belief about sabotage, it seems intuitive that $\underline{x_i}$ should be higher for males than for females. However, it turned out that the $\underline{x_i}$ for males and females were highly similar and we do not find a significant difference ($p = 0.9416$ at subject level; $p = 0.5637$ at session level).[20] This might seem puzzling but can be explained as follows: analyzing performance and the belief about the opponent's performance isolated from each other, we do not find a significant gender difference, neither for own performance nor for the beliefs about the opponent's performance. Nevertheless, males performed somewhat better (see Figure 2.1) and expected their opponent to perform on a slightly lower level than females did. Hence, males were more optimistic about their relative performance in the encoding task than females. On the other hand, they expected more sabotage from their opponents than females. The more optimistic expectation about the relative performance balanced the higher expected amount of sabotage leading to almost equal levels of $\underline{x_i}$ for both genders. Still, we observed males selecting twice as high levels of sabotage than females, and differing beliefs about the opponent's sabotage decision alone cannot explain this finding.[21] From the belief treatment we have learned that the gender gap was already

[19]Note that using the performance of the previous period as a proxy underestimates current performance due to the strong learning effects. Learning, however, does not differ between females and males (see Table A.2). Hence, there is no systematic bias against one gender when using lagged perfomance as an indicator of current performance. The results do not change if we use the performance of the respective instead of the previous period as a proxy for the expected own performance.

[20]In order to take into account that agents might learn about the opponent's behavior and therefore might form more accurate beliefs towards the end of the experiment, we also investigate the last two periods. The $\underline{x_i}$'s for females and males are also not statistically different from each other ($p = 0.4605$ at subject level; $p = 1.000$ at session level).

[21]The interpretation of the beliefs has, of course, to be treated with caution since we have point beliefs and do not know to what extent the participants considered their stated beliefs to be true.

present in the beliefs with males expecting to suffer to a much stronger degree from sabotage than females. In addition, as can be seen in Table A.5, the beliefs indeed help to explain the gender gap in sabotage to some extent. Yet, if we include beliefs about performance into the analysis (using lagged performance as a proxy for current performance), we do not find any difference between females and males in the amount of sabotage that is necessary to win the tournament.

Another possible explanation we are going to discuss is the so called "joy of winning" effect, meaning that agents derive some extra (non-monetary) utility of winning the tournament. If males derive that extra utility from winning the tournament while females do not, this might explain the gender gap we observe in sabotage. Indeed, several empirical studies suggest that winning a competition is more important to males than to females.[22] Cotton, McIntyre, and Price (2011) for instance develop a formal model (without the option to sabotage) that is consistent with previous empirical findings. Further evidence comes from Wieland and Sarin (2012). Note that males who derive some extra utility of winning might also be more likely to exert higher levels of effort to ensure they will win the tournament. While we do not find significant differences regarding performance, males tend to perform on average slightly better than females in the baseline and the belief treatment.

Concerning sabotage decisions, our results are perfectly in line with the assumption that males receive some extra utility from winning the tournament. If the status of males depends to a higher extent on outperforming opponents and feeling superior to others (Crocker, Luhtanen, Cooper, and Bouvrette (2003)), it is be important to study in what way this depends on the specific gender of the opponent.

The results of the gender treatment allow us to investigate whether agents changed their performance as well as their sabotage decisions when the gender of the opponent was revealed. In the encoding task, males, on average, achieved 121.20 points and females 106.05 points, which is a significant difference in performance ($p = 0.0136$ at subject level; $p = 0.0209$ at session level). However, we find no significant differences when we compare this performance with the performance in the baseline or the belief treatment. The data indicate that males increased their performance while females reacted less strongly to the revelation of gender, which is in line with previous findings in the literature. Given that we observe a significant gender gap in performance, one might expect that the revelation of the opponent's gender affected the gender gap in sabotage as well. On average, males selected 20.69 points of sabotage and females 10.80 points. Note that, although the change is not statistically significant, both genders chose lower average sabotage levels in the gender treatment compared to all other treatments, but the gender gap is persistent ($p = 0.0020$ at subject level; $p = 0.0209$ at session level). For a more detailed

[22] See Dohmen, Falk, Fliessbach, Sunde, and Weber (2011) for results from neuroscience that are suggestive for the existence of a general "joy of winning" effect.

analysis we have to split the sample and investigate the behavior of males and females with respect to the gender of their respective opponent separately.

Figure 2.3 gives an overview of both the performance and the beliefs about the performance of males and females when they face a male or female opponent. First, we notice that for a given gender of the opponent, males outperformed females. Second, when competing with a female, both genders performed better than when playing against a male opponent, which is consistent with previous findings in the literature (see, e.g., Price (2008), Antonovics, Arcidiacono, and Walsh (2009), Iriberri and Rey-Biel (2012), or Cotton, McIntyre, and Price (2013),). The regressions in Table A.6 in the appendix support the impression that the agent's own gender influenced the achieved points but the impact of the opponents gender is not significant if we add additional control variables. Both genders increased their performance slightly but not significantly when competing with a female opponent but the effect is prevalent for male agents.[23] The performance difference is even more striking if we incorporate the beliefs of male agents. They expected females to perform worse than male opponents but nevertheless performed better when competing with a female. Losing a competition against a female opponent seems to be worse than losing against a male competitor.

Regarding sabotage, males believed to be sabotaged more strongly than females, independent of the opponent's specific gender. This is consistent with our finding in the belief treatment. However, both genders expected a male opponent to choose a higher amount of sabotage than a female opponent. Both genders anticipated the gender gap in sabotage to some extent, but this did not lead to higher levels of sabotage when competing with a male opponent. When comparing the average belief about sabotage with the average real amount of sabotage selected by the opponent, we observe that the average difference is rather low regarding tournaments between opponents of the same gender (males vs. males 1.22 and females vs. females 1.83). In contrast, males expected females to exert more sabotage than they actually do (average difference 6.95) and females underestimate the sabotage inflicted on them by males (average difference -4.60). As can be seen in Figure 2.4 and Table A.7 in the appendix, the gender of the opponent has no significant impact on behavior, and the results are similar to those in the baseline and the belief treatment showing a significant gender gap in sabotage. Note that this gender gap is also persistent if we pool the data over all treatments and insert treatment dummies into the regressions (see Table A.8 in the appendix). It is worth mentioning that the treatment dummy for the gender treatment is significantly negative.

The remaining questions are whether it paid, in monetary terms, for males to invest in sab-

[23] If we execute the regressions for males and females separately, the gender of the opponent has a significant effect only in the regressions for males. If we further restrict our sample to agents who have competed with both genders, we have to drop 14 subjects. The results of the regressions are robust with the reduced sample, but we cannot observe a significant effect when applying a Wilcoxon matched pair signed rank test. The results can be obtained upon request.

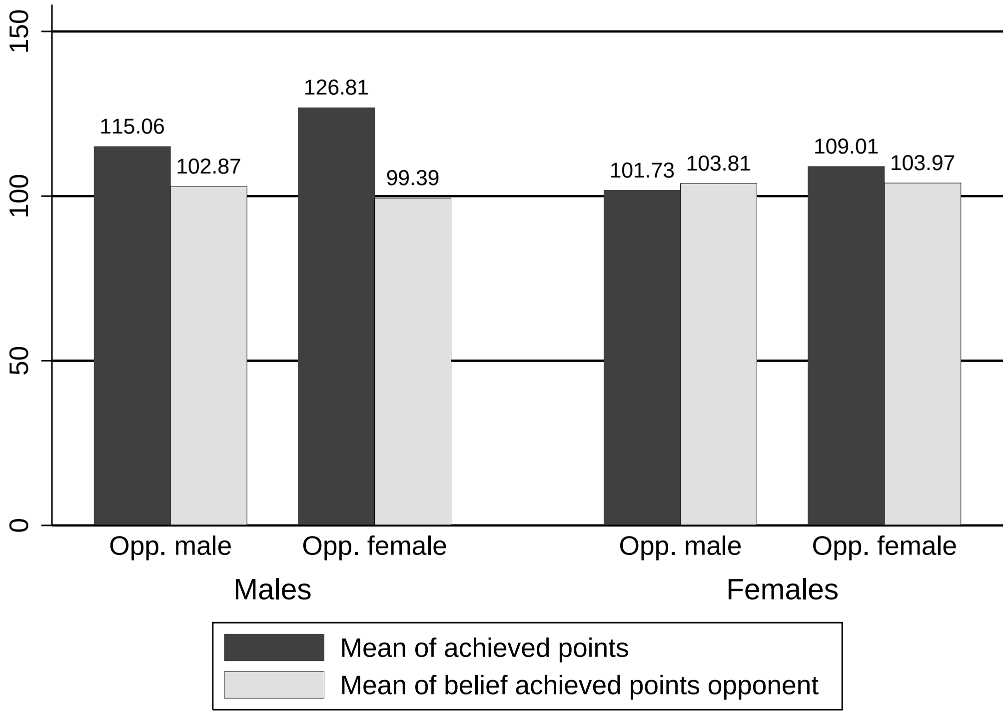

Figure 2.3.: Mean of achieved points in the gender treatment

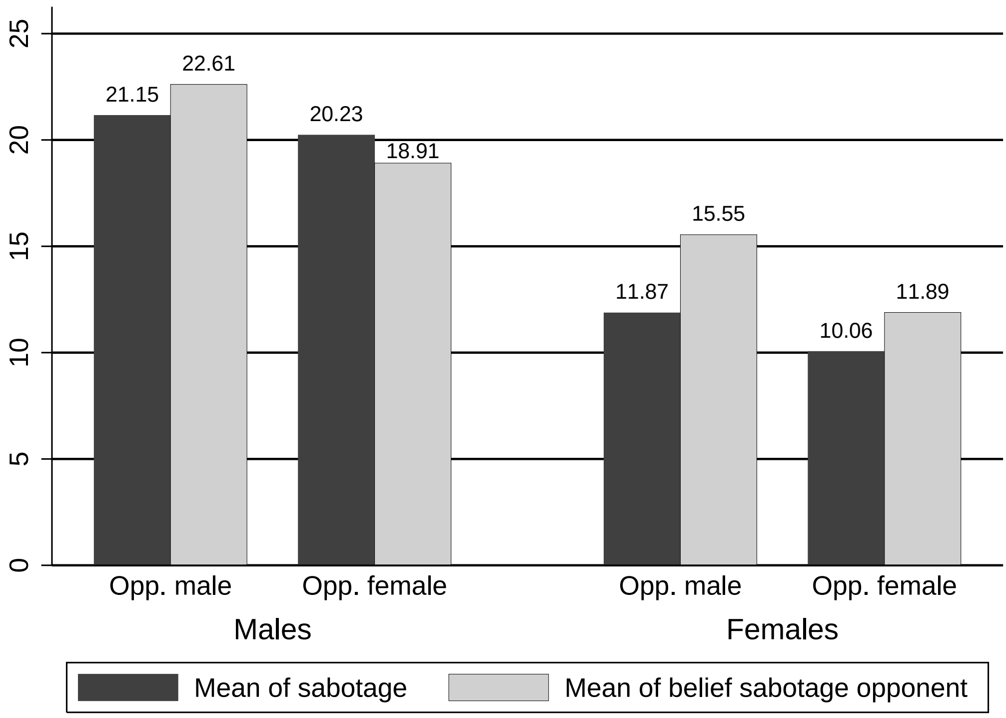

Figure 2.4.: Sabotage in the gender treatment

otage and how the gender gap affected the outcome of the tournament. Since both genders on average performed equally well in the encoding task (except for the gender treatment) and males chose a higher amount of sabotage, it is straightforward that males won the tournament more often (two-sided Fisher's exact test $p = 0.000$). Hence, males received the winner prize more frequently; but because they also had to bear higher costs the effect on earnings is ambiguous. Average earnings in the baseline and belief treatment of males were 18.28 euro, and females earned, on average, 18.22 euro. In the cheating treatment (gender treatment), male agents earned, on average, 17.77 euro (19.22 euro), whereas female agents received 18.11 euro (17.39 euro).[24] We do not find a significant gender difference concerning earnings in any treatment (baseline: $p = 0.9586$; cheating: $p = 0.4286$; belief: $p = 0.4717$; gender: $p = 0.5112$). Hence, the reluctance of females to sabotage their opponents (or to cheat) led to less success in the tournament but did not lead to less earnings compared to males. However, if males derived additional utility from winning the tournament, their overall utility was higher than that of females.

From the principal's point of view, sabotage reduced his earnings. Furthermore, the relative performance signal of the agents was biased by sabotage, and it remains an open question whether principals would expect more sabotage from males and adjust their expectations about the performance accordingly. If the outcome of the tournament was used for promotion or sorting decisions, principals systematically favored male agents. Of all male winners in the baseline treatment, about 20.40% won the tournament because of sabotage, whereas this share was 6.62% and thus much lower for females.

2.5. Conclusion

We analyze gender differences in the context of a two-player tournament with real-effort and the possibility to manipulate output. Males and females systematically differed in their sabotage decisions, and males selected on average twice as much sabotage as females. The gender gap in sabotage is large and highly significant in all treatments. As for the performance in the encoding task, we find no significant difference in the number of achieved points when the gender of the respective opponent was not revealed.

Due to the higher amount of sabotage by males, they won the tournament more often. But higher amounts of sabotage also led to higher costs, and therefore average earnings did not differ significantly between genders. Our data revealed that the gender gap was already present in the beliefs of the agents. Social preferences with respect to the principal, risk aversion, or

[24] In order to make earnings comparable between treatments, we adjusted the earnings from the belief and gender treatment for the incentivized belief elicitation.

human values, such as for instance the importance of power or benevolence, do not fully explain the differences in sabotage.

Based on previous findings and the data of the belief and the gender treatment, we believe it is likely that males derived extra utility from winning the competition and were therefore willing to invest money to ensure their victory. Females, on the other hand, are usually described as less status seeking. They were, therefore, not willing to invest that much in sabotage. However, they were aware that, on average, their opponents would choose higher levels of sabotage and that males would, on average, choose higher levels of sabotage than female opponents.

Our results show that in the encoding task, the disclosure of the opponent's gender led to a higher performance of males compared to females while leaving sabotage almost unchanged. This raised the total output and this effect was strongest when the opponent was female. Hence, in our experiment a principal achieved higher output levels in mixed tournaments where the gender of the opponent is revealed. Of course, in organizations managers should try to prevent sabotage as it leads to distorted outcomes. Although this recommendation is known in the literature, we provide an additional reason why sabotage might be harmful for organizations because it might lead to systematically wrong promotion decisions. In our data about 20% of the male winners achieved their victory based on sabotage rather than performance. Hence, 20% of all male winners would probably be granted a promotion although they were not the better performers in the encoding task.

Companies that have established gender quotas, or consider implementing affirmative action programs, need to take into account that sabotage might bias the results of tournaments and that the signals might be biased at the expense of females. The reluctance of females to compete or apply for jobs for which they are well qualified might also be affected by a fear of being sabotaged. Sabotage reduces the chance to win, and the refusal to enter the tournament might therefore be a rational decision.

3. Expectation-Based Loss Aversion and Strategic Interaction

3.1. Introduction

Next to Expected Utility Theory, Kahneman and Tversky's Prospect Theory (1979) has become the most prominent approach for modeling risk preferences. Beside probability weighting, the central building blocks of Prospect Theory are reference dependence and loss aversion—i.e., every outcome is coded as a gain or a loss relative to some value-neutral reference point and losses loom larger than equally sized gains. In a series of papers, Kőszegi and Rabin (2006, 2007, 2009) propose a theoretical framework of how a decision maker's reference point is shaped by rational expectations.[1,2] In individual decision contexts, their model has been fruitfully applied to explain a wide range of phenomena that are hard to reconcile with the standard notion of risk aversion—e.g., often observed price stickiness (Heidhues and Kőszegi, 2008), the prevalence of flat-rate tariffs (Herweg and Mierendorff, 2013), or the widespread use of bonus contracts (Herweg, Müller, and Weinschenk, 2010). Without doubt, however, many economically relevant outcomes are not determined by isolated individual decision making but by the interplay of several individuals who interact strategically. The few contributions that analyze strategic interaction of expectation-based loss-averse players do so in rather specific environments—e.g., rank-order tournaments (Gill and Stone, 2010), auctions (Lange and Ratan, 2010), team production (Daido and Murooka, 2014). Moreover, these contributions do not consider the possibility of mixed strategy equilibria and often even restrict attention to specific sets of pure strategy equilibria, e.g., symmetric equilibria. Thus, up to date, we lack a general understanding of the

[1]The general feature that the reference point is shaped by forward-looking expectations is shared with the disappointment aversion models of Bell (1985), Loomes and Sugden (1986), and Gul (1991). In the remainder of the paper, however, whenever we speak of (expectation-based) loss aversion, we do so in the sense of Kőszegi and Rabin.

[2]Empirical evidence supporting the theory of Kőszegi and Rabin is provided by Abeler, Falk, Goette, and Huffman (2011), Crawford and Meng (2011), Ericson and Fuster (2011), and Gill and Prowse (2012).

overarching patterns how expectation-based loss aversion affects players' strategic interaction.

In this paper, we provide a comprehensive analysis regarding strategic interaction under expectation-based loss aversion. The resulting insights correspond to the following contributions: First, we develop a coherent analytical framework by extending the equilibrium concepts of Kőszegi and Rabin (2006, 2007) to finite games and explain the methodology how to derive such equilibria. Second, we identify three major characteristics of the strategic behavior of expectation-based loss-averse agents that differ from the behavior of agents with standard expected-utility preferences: *decisiveness* and *adaptiveness* for fixed expectations, and *reluctance to mix* for choice-acclimating expectations. Third, Third, we analyze equilibrium play under expectation-based loss aversion and address the question of equilibrium existence.

Kőszegi and Rabin (2006) focus on situations where the decision maker ponders a future decision and forms expectations about her actions before she actually takes action. In these situations a *personal equilibrium (PE)*, essentially, requires internal consistency, i.e., only to make plans that one is willing to follow through later on. We define a *personal Nash equilibrium (PNE)* as a strategy profile such that each player plays a PE given her opponents' behavior. Complementary, Kőszegi and Rabin (2007) consider situations where the decision maker is confronted with the decision to be made rather unexpectedly. In this case, the action taken necessarily coincides with the decision maker's plan. The choice of the most desirable course of action is referred to as the *choice-acclimating personal equilibrium (CPE)*. We define a *choice-acclimating personal Nash equilibrium (CPNE)* as a strategy profile such that all players play a CPE given the opponents' behavior.

Expectation-based loss aversion represents an alternative to Expected Utility Theory for modeling risk preferences. When focusing on pure strategies in games without inherent uncertainty, the game is devoid of risk. As a consequence, we find that equilibrium predictions are identical under Nash equilibrium, PNE, and CPNE. Once the consequences of players' actions become risky, this picture changes significantly. If any player plays a mixed strategy or there is a draw of nature, then the derivation of equilibria, best-response behavior, and equilibrium play differ for expectation-based loss-averse players in comparison to their counterparts with standard expected-utility preferences.

The derivation of (mixed) Nash equilibria for players with standard expected-utility preferences relies upon the fact that a player's expected utility is linear in each of the probabilities that she attaches to her own pure strategies. In consequence, if a player with standard preferences is willing to play some particular probabilistic mixture over a given set of pure strategies, she is willing to play any (possibly degenerate) mixture over this set of pure strategies. Furthermore, if her opponents change their behavior slightly, she typically will not be willing to mix over the same set of pure strategies anymore. In light of these observations, mixed strategy equilibria under Expected Utility Theory have been controversially discussed and are regarded as intuitively problematic.[3]

We identify three behavioral features of expectation-based loss-averse players which set their strategic behavior distinctively apart from players with expected-utility preferences. A loss-averse player's expected utility from playing a particular pure strategy depends on her expectations regarding her own behavior. Hence, the attractiveness of a pure strategy can only be assessed for a *given plan of action*. If a player's plan assigns rather high (low) probability to a specific pure strategy, she becomes attached to the idea that the associated outcomes will (not) occur. Due to this attachment, the player then may actually prefer to play this strategy with certainty (not at all). Either way, she is not willing to stick to her original plan. We find that there exists at most one plan of action which balances such diverging attachments and makes different pure strategies equally attractive. We refer to this behavioral feature as *decisiveness*, because there exists at most one mixed PE over a given set of pure strategies. The second distinguishing feature of the strategic behavior of expectation-based loss-averse players is *adaptiveness*: if a loss-averse player is willing to mix over a given set of pure strategies and her opponents' strategies change slightly, she remains willing to mix over the very same set of pure strategies irrespective of the exact behavioral change. Arbitrary changes of the opponents' strategies lead to a change in a player's expected material utility induced by any of her pure strategies. Since expectations directly influence her utility, however, there always exists a slight adaption in expectations that exactly counteracts this change in material utility. Hence, the player is willing to follow through the adapted plan such that slight arbitrary trembles in her opponents' behavior do not wipe out her willingness to mix over the same set of pure strategies. Thus, the concept of a mixed strategy is—in a very literal sense—more robust under loss aversion with fixed expectations than under standard expected-utility preferences.

For the case of choice-acclimating expectations, in contrast, loss-averse players exhibit a general *reluctance to mix*. The reason is that a loss-averse player with choice-acclimating expectations strongly desires to reduce risk, which she can achieve by choosing a pure strategy rather than a mixture between several pure strategies. Therefore, a mixture over several pure strategies decreases her expected utility even if she is indifferent between these. Consequently, behavior compatible with choice-acclimating expectations never involves mixing over several pure strategies if the probabilistic consequences of these pure strategies are not identical.

Finally, the characteristics of the strategic behavior of expectation-based loss-averse players have direct implications for equilibrium play and existence. Since players with fixed expectations are decisive, a player's PE correspondence is not necessarily convex valued. In consequence, Kakutani's fixed point theorem is not applicable and the existence of a PNE is a priori unclear. For two-player games with two pure strategies for each player, however, we show that adaptiveness induces the graph of a player's PE correspondence to be connected. Hence,

[3]For surveys regarding the interpretation of mixed strategies see Aumann (1985) and Rubinstein (1991).

in this basic case, a PNE always exists. Furthermore, we show that expecting to play a materially weakly dominant strategy always constitutes a credible plan. Therefore, whenever a game features a Nash equilibrium in weakly dominant strategies, existence of a PNE is ensured. Also, expecting to play any other strategy is not a credible plan. Hence, if there exists a Nash equilibrium in materially weakly dominant strategies, this constitutes the unique PNE.

For choice-acclimating beliefs the step from players' CPE correspondences to CPNE is even more apparent. As players are reluctant to mix over pure strategies in this case, a CPNE can never involve mixed strategies. Hence, the existence of a CPNE is not guaranteed. More specifically, we show that existence of CPNE can fail even in basic games without inherent uncertainty.[4] This insight raises the question if there are conditions that guarantee the existence of a CPNE. We show that a Nash equilibrium in weakly dominant strategies always constitutes the unique CPNE of the game, which implies existence for this case. Hence, in public good games a CPNE always exists—even if there is uncertainty about the other players' endowment. More specifically, the tendency to free ride and not to contribute remains an equilibrium under loss aversion. Similarly, in the Vickrey auction it is a CPNE to bid the true valuation. On the one hand, the potential non-existence calls into question how suited CPNE is for the analysis of strategic interaction. On the other hand, the absence of mixed strategy CPNEs complements existing and future contributions that study strategic interaction of expectation-based loss-averse players on the basis of pure strategy equilibria in applications like auctions, rank-order tournaments, or team production. They can rest assured that a focus on pure strategy CPNEs is without loss of generality.

The rest of the paper is organized as follows. Section 3.2 provides a brief overview over the theoretical literature that applies expectation-based loss aversion à la Kőszegi and Rabin both to individual decision making and strategic interaction. Section 3.3 formally introduces the class of games we study while Section 3.4 extends the equilibrium concepts PE and CPE to strategic interaction. In Section 3.5, we demonstrate the derivation of PEs and CPEs in situations of strategic interaction and analyze the resulting behavior of expectation-based loss-averse players. Section 3.6 comments on equilibrium play under expectation-based loss aversion and the existence of PNEs and CPNEs. We provide a discussion of alternative interpretations of mixed strategies and multidimensional outcomes in Section 3.7. Section 3.8 concludes.

[4] As we will lay out in more detail, the potential non-existence of CPNE is rooted in the notion that each player individually randomizes over the set of her pure strategies. Under the interpretations of mixed strategies according to Rosenthal (1979) or Aumann and Brandenburger (1995) a CPNE always exists.

3.2. Related Literature

By now, a plethora of theoretical contributions analyzes individual (i.e., nonstrategic) decision making in a variety of economic environments when agents are expectation-based loss averse à la Kőszegi and Rabin. One strand of research considers risk- and loss-neutral firms selling to expectation-based loss-averse consumers. Here, Heidhues and Kőszegi (2008) show how consumer loss aversion can account for focal pricing, i.e., nonidentical competitors charging identical prices for differentiated products.[5] Herweg and Mierendorff (2013) find that uncertainty about their own future demand leads to consumers preferring a flat rate to a measured tariff, which in turn can make the profit-maximizing contract to be offered by firms a flat rate. Analyzing product-availability strategies, Rosato (2014b) shows that limited-availability sales can manipulate consumers into an ex-ante unfavorable purchase by raising the consumers' reference point through a tempting discount on a good available only in limited supply.[6] Another strand analyzes optimal incentive provision with expectation-based loss-averse agents. Herweg, Müller, and Weinschenk (2010) show that the optimal incentive contract takes the form of a simple binary payment scheme even if the performance measure is arbitrarily rich. Applying the dynamic loss-aversion model by Kőszegi and Rabin (2009), Macera (2013) studies the intertemporal allocation of incentives in a repeated moral hazard model.[7] Furthermore, the concept of expectation-based loss aversion à la Kőszegi and Rabin has been applied to questions of inventory management (Herweg, 2013), task assignment (Daido, Morita, Murooka, and Ogawa, 2013), and incomplete contracting (Herweg, Karle, and Müller, 2014).

Recently, a number of contributions began to address strategic interaction of expectation-based loss-averse individuals in rather specific environments of economic interest. In the context of rank-order tournaments, Gill and Stone (2010) show that even with symmetric contestants the only stable CPNEs are asymmetric if loss aversion is sufficiently important. Analyzing the optimal structure of team compensation, Daido and Murooka (2014) find that the optimal wage scheme can display team incentives even when individual success probabilities are independent because this reduces the agents' expected losses. Particular interest has been drawn to the behavior of expectation-based loss-averse bidders in auctions. Lange and Ratan (2010) use CPNE as a solution concept for first- and second-price sealed-bid auctions, showing that expectation-based loss aversion can explain overbidding relative to the Nash prediction

[5] Karle and Peitz (2014) study the implications for competitiveness of the market outcome if some consumers are initially uninformed about their tastes and form a reference point consisting of an expected match-value and price distribution. Considering a monopolistic seller, Heidhues and Kőszegi (2014) explain the occurrence of sales.

[6] Karle (2014) analyzes how a monopolist can manipulate consumers' willingness to pay by disclosing verifiable product information.

[7] Daido and Itoh (2007) study self-fulfilling prophecies in the form of the Galatea and the Pygmalion effect.

in induced-value auctions. Extending this work, Belica and Ehrhart (2014) consider how the results change if PNE is applied. Eisenhuth (2010) demonstrates that for loss-averse bidders with choice acclimating beliefs, the revenue-maximizing auction is an all pay auction with minimum bid. All of these papers investigate auctions that have only one period. Analyzing sequential two-round sealed-bid auctions, Rosato (2014a) shows that prices of identical goods tend to decline between rounds in a sequential CPNE, i.e., expectations-based loss aversion can rationalize the empirically well-documented "afternoon-effect". Applying PNE, Ehrhart and Ott (2014) show the differences in behavior of loss-averse bidders between English and Dutch auctions.

Closest in spirit to our paper is Shalev (2000), who also analyzes strategic interaction of loss-averse individuals. Regarding reference point formation, however, he follows Gul (1991) and assumes that the reference point corresponds to a lottery's certainty equivalent in utility terms given that the lottery is evaluated with respect to that reference point. In consequence, the reference point is not a lottery over outcomes—as in Kőszegi and Rabin (2006, 2007)—but a single point. Under this concept of reference point formation, Shalev (2000) gives a general account of equilibrium existence and compares pure strategy Nash equilibria to equilibria played by loss-averse players for games with perfect information. Due to the different approaches how expectations shape a player's reference point, however, the strategic deliberations of loss-averse players identified by Shalev (2000) are rather different from those identified in this paper—most notably, they are neither decisive, nor adaptive, nor reluctant. In consequence, with considerations of loss-averse players regarding the use of mixed strategies resembleing those of players with standard preferences, equilibrium existence in Shalev (2000) is guaranteed by Kakutani's fixed point theorem.

3.3. The Model

For the analysis of strategic interaction between expectation-based loss-averse players we consider finite games with the following elements. First, the set of players denoted by $\mathcal{I} = \{1, \ldots, I\}$ is finite. Second, each player $i \in \mathcal{I}$ has a finite pure-strategy space $\mathcal{S}^i = \{s^i_1, \ldots, s^i_{M^i}\}$. A pure-strategy profile is denoted by $s = (s^1, \ldots, s^I) \in \mathcal{S}$, where $\mathcal{S} = \times_{i=1}^{I} \mathcal{S}^i$. Third, there is a finite set $\Theta = \{\theta_1, \ldots, \theta_N\}$, where the elements of Θ are realizations of some random variable θ which is determined by a draw of nature. We denote the probability of θ_j being drawn by $Q(\theta_j) \geq 0$. Fourth, each player $i \in \mathcal{I}$ has payoff function $u^i : \mathcal{S} \times \Theta \to \mathcal{U}^i \subset \mathbb{R}$ which maps any combination of a pure-strategy profile $s \in \mathcal{S}$ and randomly drawn $\theta_j \in \Theta$ into a material payoff $u^i(s, \theta_j) \in \mathbb{R}$.

In this setting, a mixed strategy $\sigma^i = (\sigma^i(s^i_1), \ldots, \sigma^i(s^i_{M^i}))$ for player $i \in \mathcal{I}$ is a lottery over her pure strategies, where $\sigma^i(s^i_m)$ denotes the probability of player i playing the pure strategy

s^i_m. The space of player i's mixed strategies is denoted by Σ^i. Accordingly, the space of mixed-strategy profiles $\sigma = (\sigma^1, \ldots, \sigma^I)$ is $\Sigma = \times_{i=1}^I \Sigma^i$. As usual, we will sometimes refer to the mixed strategy profile σ as (σ^i, σ^{-i}), where $\sigma^{-i} \in \Sigma^{-i} = \times_{j \neq i} \Sigma^j$ denotes the mixed-strategy profile for all players except player i.[8]

We assume players to be loss averse à la Kőszegi and Rabin (2006). Hence, the overall utility that player $i \in \mathcal{I}$ derives from some riskless material payoff u consists of two components: traditional material utility given by u itself and psychological gain-loss utility. Gain-loss utility is determined by a comparison of the material payoff u to some reference material payoff u^r. The player feels a gain if the payoff u exceeds the reference payoff u^r, otherwise she suffers a loss. Formally, overall gain-loss utility is given by $\mu(u - u^r)$, where $\mu(\cdot)$ denotes the so-called value function according to which the deviation from the reference outcome is evaluated. We assume the value function to be piece-wise linear:

$$\mu(u - u^r) = \begin{cases} \eta(u - u^r) & \text{if} \quad u \geq u^r \\ \eta\lambda(u - u^r) & \text{if} \quad u < u^r \end{cases} . \tag{3.1}$$

Here, $\eta \geq 0$ denotes the weight the player puts on psychological gain-loss utility relative to intrinsic material utility and $\lambda > 1$ captures loss aversion, i.e., losses loom larger than gains of equal size.[9]

A player's reference point corresponds to a reference lottery over her potential material payoffs which is determined by her expectations about her own strategy and the strategies played by the other players. Let $\Lambda^i(u) = \{(s, \theta) \in S \times \Theta \,|\, u^i(s, \theta) = u\}$ denote the set of (s, θ) combinations that result in some specific material payoff $u \in \mathcal{U}^i$ for player $i \in \mathcal{I}$. The probability of this payoff for player $i \in \mathcal{I}$ being realized under strategy profile σ is given by $P^i(u|\sigma) = \sum_{(s,\theta) \in \Lambda^i(u)} Q(\theta) \Pi_{j=1}^I \sigma^j(s^j)$. Hence, if player i expects the opponents to play σ^{-i} and herself to play $\hat{\sigma}^i$, she expects payoff $u \in \mathcal{U}^i$ to be realized with probability $P^i(u|(\hat{\sigma}^i, \sigma^{-i}))$. Given these expectations, her overall expected utility from playing strategy profile σ^i is given by

$$\begin{aligned} U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i}) = & \sum_{u \in \mathcal{U}^i} P^i(u|(\sigma^i, \sigma^{-i})) \cdot u \\ & + \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|(\sigma^i, \sigma^{-i})) \cdot P^i(\tilde{u}|(\hat{\sigma}^i, \sigma^{-i})) \cdot \mu(u - \tilde{u}). \end{aligned} \tag{3.2}$$

The first part of overall expected utility reflects expected material utility, where the expectation is based on the lottery over feasible material payoffs induced by strategy σ^i that player i actually

[8] Clearly, σ^i might also be a degenerate lottery and thus represent a pure strategy. If we want to be explicit about player i playing a pure strategy, however, we usually write the strategy profile σ as (s^i, σ^{-i}).

[9] Almost all of the contributions cited in Section 3.2 use this piece-wise linear specification of the gain-loss function.

plays. The second part reflects expected gain-loss utility, where each material payoff that could possibly be realized is compared with every other feasible material payoff. Here, each such comparison is weighted by its occurrence probability based on the expectation $\hat{\sigma}^i$ that player i holds with regard to her own strategy.[10]

3.4. Equilibrium Concepts

For the context of individual decision making, Kőszegi and Rabin propose two different notions of equilibrium for consistent behavior of expectation-based loss-averse individuals. These two notions differ with regard to the timing when expectations about the decision in question are formed and when this decision is actually taken.

Personal equilibrium (PE) applies to situations where a person has some time to ponder about a decision before she is called to make her choice. Here, with the person thinking about—but not being able to commit to—her choice before making it, she will enter the actual decision with previously formed and thus fixed expectations regarding her own behavior. At the moment of choice, however, the individual might prefer to deviate from what she expected to do—maybe because she relishes the idea of saving some money or effort cost which she originally planned to invest or to exert, respectively. In this case, the individual should have foreseen that her course of action will not meet her expectations, such that she should not have expected to act this way in the first place. Therefore, PE requires internal consistency in the sense that a person can reasonably expect a particular course of action only if she is willing to follow it through given her expectations. The following definition extends this idea to a situation of strategic interaction, where all players have some time to ponder their own behavior before choosing their strategy of play.

Definition 3.1. *A personal Nash equilibrium (PNE) is a vector $\sigma \in \Sigma$ such that for each player $i \in \mathcal{I}$,*

$$U^i(\sigma^i, \sigma^i, \sigma^{-i}) \geq U^i(\tilde{\sigma}^i, \sigma^i, \sigma^{-i}), \quad \forall\, \tilde{\sigma}^i \in \Sigma^i$$

According to Definition 3.1, a PNE is a vector of (possibly mixed) strategies such that each player is willing to follow through with the strategy she expected to play given the other players' strategies. Thus, in a PNE, every player plays a PE in the sense of Kőszegi and Rabin (2006).

Choice-acclimating personal equilibrium (CPE) addresses situations where a person is called to make her choice without having much time to contemplate this choice. In this case, the person's expectations are not fixed but literally dictated by her behavior. Hence, she expects exactly those consequences that actually prevail given her chosen course of action. With no

[10] With regard to sequential games, we abstract from players updating their expectations as play proceeds, i.e., we do not allow for players experiencing paper gains or paper losses as considered in Kőszegi and Rabin (2009).

		player 2	
		go straight	swerve
player 1	go straight	0,0	3,1
	swerve	1,3	2,2

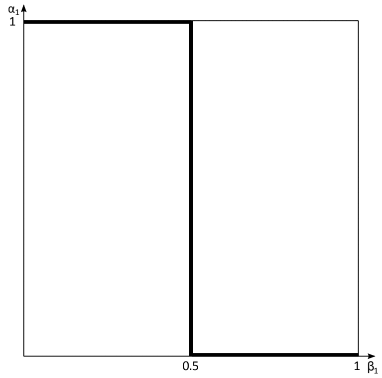

Figure 3.1.: Material utility payoff matrix and best response curve for player with standard expected utility preferences.

scope for expectations to diverge from actual behavior, CPE requires internal consistency in the sense of the person taking the course of action that maximizes her expected well-being. Definition 3.2 extends this idea to a situation of strategic interaction, where all players have to choose their strategy of play without having time to ponder their own behavior in advance.

Definition 3.2. *A choice-acclimating personal Nash equilibrium (CPNE) is a vector $\sigma \in \Sigma$ such that for each player $i \in \mathcal{I}$,*

$$U^i(\sigma^i, \sigma^i, \sigma^{-i}) \geq U^i(\tilde{\sigma}^i, \tilde{\sigma}^i, \sigma^{-i}), \quad \forall\, \tilde{\sigma}^i \in \Sigma^i.$$

According to Definition 3.2, a CPNE is a vector of (possibly mixed) strategies such that each player chooses a strategy—and at the same time adopts her expectations about possible future outcomes according to that choice—that maximizes her expected utility given the other players' strategies. Hence, in a CPNE, every player plays a CPE in the sense of Kőszegi and Rabin (2007).[11]

3.5. Strategic Behavior of Loss-averse players

In this section, we derive general insights how the reasoning behind the strategic behavior and its derivation differs for expectation-based loss-averse players compared to their counterparts with standard expected-utility preferences. In order to convey these insights and their intuition more vividly, we will repeatedly refer to the simple example of the anti-coordination game known as "Chicken", which is depicted in Figure 3.1.

[11] According Definitions 3.1 and 3.2 all players form their expectations at the same point in time; i.e., all players' expectations are either fixed or choice acclimating. Amending the above concepts to allow for situations where some players have fixed expectations while other players have choice-acclimating expectations is straightforward. In particular, all results in Section 3.5 remain valid under such a modification.

The story of the Chicken game is well known: Two drivers head for a single-lane bridge from opposite directions and each player has to decide whether she goes straight for the bridge or swerves. The first driver to swerve away yields the bridge to the opponent. While her opponent thereafter will brag about her victory and be celebrated as a daredevil, a man without fear, the driver who swerved will be publicly regarded as a coward. If both players swerve, there is nothing to brag about and each driver has to live with the silent shame of having chickened out. Finally, if neither player swerves, the result is a close-to-fatal crash in the middle of the bridge. As is reflected in the material utility values in Figure 3.1, the best possible outcome is to be the public hero and the worst possible outcome is to be in a severe car crash. Furthermore, a life in public shame is worse than a life in silent shame.

The pure strategy space for player $i = 1, 2$ is $\mathcal{S}^i = \{\text{go straight}, \text{swerve}\}$. To ease notation, we denote $\sigma^1(\text{go straight}) = \alpha_1$ and $\sigma^2(\text{go straight}) = \beta_1$, where $0 \leq \alpha_1, \beta_1 \leq 1$. Likewise, $\sigma^1(\text{swerve}) = \alpha_2$ and $\sigma^2(\text{swerve}) = \beta_2$, where $\alpha_2 = 1 - \alpha_1$ and $\beta_2 = 1 - \beta_1$. If one driver is more likely to go straight (swerve), the other driver maximizes her expected material utility by swerving (going straight). Only if one driver goes straight with the same probability as she swerves, the expected material utility from swerving equals the expected material utility from going straight for the other player, making her indifferent between going straight and swerving. These observations are summarized in the right panel of Figure 3.1, which depicts player 1's best response in terms of the optimal probability α_1 to go straight for a given probability β_1 to go straight of player 2.

3.5.1. Redundancy of Pure Strategies

To pave the way for the following analysis of strategic interaction, we next introduce the definition of a *redundant pure strategy*. To this end, let

$$L^i(\sigma^i, \sigma^{-i}) = \left(P(u|\sigma^i, \sigma^{-i})\right)_{u \in \mathcal{U}^i} \tag{3.3}$$

denote the lottery over the set of material utility outcomes for player i which is induced by player i playing strategy σ^i and her opponents playing the strategy profile σ^{-i}.

Definition 3.3. *Given her opponents' strategy profile σ^{-i}, player i's pure strategy s^i_k is redundant if and only if there exists a set of pure strategies $\tilde{\mathcal{S}}^i \subseteq \mathcal{S}^i \setminus \{s^i_k\}$ and numbers $(\gamma(s^i))_{s^i \in \tilde{\mathcal{S}}^i}$ such that $L^i(s^i_k, \sigma^{-i}) = \sum_{s^i \in \tilde{\mathcal{S}}^i} \gamma(s^i) L^i(s^i, \sigma^{-i})$.*

A pure strategy s^i_k is redundant if the lottery over material utility outcomes induced by s^i_k is a linear combination of the lotteries induced by a set of pure strategies not containing s^i_k. Note that pure strategy s^i_k being non-redundant implies that player i cannot replicate the lottery over outcomes induced by s^i_k by playing any mixed strategy excluding s^i_k. In this sense, only a non-redundant strategy bears importance for player i as its probabilistic outcome consequences are unique. Note that the pure strategies in our leading example of the Chicken game are never redundant, because going straight leads to a different material utility outcome than swerving for every possible strategy σ^2 of the other driver.[12]

The following lemma states a very helpful observation which we will evoke repeatedly in the formal analysis of strategic behavior of expectation-based loss-averse players.

Lemma 3.1. *Let $L^j = (p^j(u))_{u\in\mathcal{U}}$ with $j \in \{A, B\}$ denote two lotteries over some finite outcome space $\mathcal{U} \subset \mathbb{R}$, where $p^j(u)$ denotes the probability that outcome $u \in \mathcal{U}$ is realized under L^j. Then*

$$\frac{\sum_{u\in\mathcal{U}}\sum_{\tilde{u}\in\mathcal{U}} p^A(u)p^A(\tilde{u})|u-\tilde{u}| + \sum_{u\in\mathcal{U}}\sum_{\tilde{u}\in\mathcal{U}} p^B(u)p^B(\tilde{u})|u-\tilde{u}|}{2} \leq \sum_{u\in\mathcal{U}}\sum_{\tilde{u}\in\mathcal{U}} p^A(u)p^B(\tilde{u})|u-\tilde{u}|, \quad (3.4)$$

with (3.4) *holding with equality if and only if L^A and L^B are identical.*

Lemma 3.1 states that the expected difference between two draws from different lotteries L^A and L^B is larger than the average of the expected difference between two draws from lottery L^A and the expected difference between two draws from lottery L^B. Note how Lemma 3.1 relates to the above definition of redundancy: If two pure strategies s^i_k and s^i_m induce lotteries $L^i(s^i_k, \sigma^{-i})$ and $L^i(s^i_m, \sigma^{-i})$ for which (3.4) holds with equality, these lotteries are identical and the pure strategies s^i_k and s^i_m are redundant.

Lemma 3.1 has important implications for an expectation-based loss-averse player's inclination to play mixed strategies. Kőszegi and Rabin (2007) refer to $\sum_{u\in\mathcal{U}}\sum_{\tilde{u}\in\mathcal{U}} p^j(u)\, p^j(\tilde{u})|u-\tilde{u}|$ as the average self-distance of lottery $L^j = (p^j(u))_{u\in\mathcal{U}}$ with finite support $\mathcal{U}$.[13] The average self-distance of lottery L^j is inversely proportional to the psychological gain-loss utility associated with playing (and expecting to play) L^j—*inversely* because each comparison effectively enters as a net loss due to loss aversion. Thus, the average-self distance is a measure for the psychological disutility arising from being exposed to the riskiness embodied in lottery L^j. The gain-loss utility from playing (and expecting to play) a probabilistic mixture of two lotteries L^A and L^B, on the other hand, comprises not only within-lottery comparisons, but also comparisons across lotteries, where each comparison again enters as a net loss. In consequence, the gain-loss utility associated with randomizing between lotteries L^A and L^B decreases not only in the average-self distances of L^A and L^B, but also in the average distance between the lotteries L^A and L^B, which

[12]To give an example for redundant strategies in a simple two-by-two game, consider a symmetric version of *Matching Pennies*. If player 2 plays heads and tails with equal probability, the probabilistic outcome consequences of both pure strategies are identical for player 1 and the pure strategies are redundant.

[13]See Definition 5 on p.1063 in Kőszegi and Rabin (2007) for the definition of a lottery's average self-distance.

is given by $\sum_{u \in \mathcal{U}} \sum_{\tilde{u} \in \mathcal{U}} p^A(u) p^B(\tilde{u}) |u - \tilde{u}|$. According to Lemma 3.1, however, the latter exceeds the average of the average self-distances, such that either the gain-loss utility associated with L^A or the gain-loss utility associated with L^B is less negative than the gain-loss utility associated with any randomization over lotteries L^A and L^B. In this sense, randomizing over lotteries creates an additional layer of uncertainty that a loss-averse player in tendency dislikes.

This observation suggests that the willingness to play a mixed strategy for an expectation-based loss-averse player should be limited in comparison to a player with standard expected-utility preferences. In the remainder of this section, we establish that this conjecture holds for both fixed expectations and choice-acclimating expectations—albeit to a differing degree.

3.5.2. Personal Equilibrium

We start by deriving the set of PEs for player i given her expectations regarding her own behavior and her opponents' strategies. Recall that player i's expected utility $U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$ is linear in each component $\sigma^i(s^i_m)$ of player i's strategy σ^i. Hence, the marginal utility of $U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$ with respect to $\sigma^i(s^i_m)$ does not depend on any component of player i's strategy σ^i. In other words, given the strategy profile of player i's opponents, σ^{-i}, and her expectations regarding her own choice of strategy, $\hat{\sigma}^i$, moving probability mass from one pure strategy s^i_m to some other pure strategy s^i_k changes expected utility at a constant rate equal to $\partial U^i / \partial \sigma^i(s^i_k) - \partial U^i / \partial \sigma^i(s^i_m)$. Consequently, marginal expected utilities reflect the attractiveness of the associated pure strategies given σ^{-i} and $\hat{\sigma}^i$.

Marginal expected utilities are then suitable to characterize the set of PEs. Given her opponents' strategies σ^{-i}, a pure strategy s^i_m is a PE for player i if and only if the marginal expected utility with respect to component $\sigma^i(s^i_m)$ is among the greatest marginal expected utilities given $\hat{\sigma}^i(s^i_m) = 1$ and $\hat{\sigma}^i(s^i_k) = 0$ for all $k \neq m$. Hence, s^i_m is at least as attractive as any other pure strategy since moving probability mass from s^i_m to any other pure strategy weakly decreases expected utility. Similarly, a mixed strategy σ^i is a PE if, given player i expects to play σ^i, she does not prefer to depart from this plan. Let

$$\Gamma(\sigma^i) = \{s^i_m \in \mathcal{S}^i \mid \sigma^i(s^i_m) > 0\} \tag{3.5}$$

denote the set of pure strategies which are played with strictly positive probability under strategy σ^i. The cardinality of $\Gamma(\sigma^i)$, which is denoted by $|\Gamma(\sigma^i)|$, then specifies the number of pure strategies that are played with positive probability under σ^i. Mixed strategy σ^i constitutes a PE if, given $\hat{\sigma}^i = \sigma^i$, all pure strategies in $\Gamma(\sigma^i)$ are equally attractive for the player and at least as attractive as all other strategies. Formally, this is the case if and only if for any $s, s' \in \Gamma(\sigma^i)$ and $s'' \notin \Gamma(\sigma^i)$ we have $\partial U^i / \partial \sigma^i(s) = \partial U^i / \partial \sigma^i(s') \geq \partial U^i / \partial \sigma^i(s'')$.

Denote by

$$R^i(\sigma^{-i}) = \{\sigma^i \in \Sigma^i \mid U^i(\sigma^i, \sigma^i, \sigma^{-i}) \geq U^i(\tilde{\sigma}^i, \sigma^i, \sigma^{-i}) \ , \ \forall \, \tilde{\sigma}^i \in \Sigma^i\} \tag{3.6}$$

the correspondence of PEs for player i. Proposition 3.1 summarizes two important features of the PE correspondence which qualitatively set it apart from a best response correspondence of a player with standard expected-utility preferences.[14]

Proposition 3.1. *Suppose $\bar{\sigma}^i \in \Sigma^i$ is a mixed strategy PE for player i given the strategy profile σ^{-i} of her opponents, i.e., $\bar{\sigma}^i \in R(\sigma^{-i})$ with $|\Gamma(\sigma^i)| \geq 2$, that does not involve any redundant strategies.*

(i) ***Decisiveness:*** *There is no other mixed strategy PE for player i that involves mixing over the same pure strategies as $\bar{\sigma}^i$:*

$$\nexists \tilde{\sigma}^i \in R(\sigma^{-i}) \ \text{ s.t. } \ \tilde{\sigma}^i \neq \bar{\sigma}^i \text{ and } \Gamma(\tilde{\sigma}^i) = \Gamma(\bar{\sigma}^i).$$

(ii) ***Adaptiveness:*** *Suppose $\bar{\sigma}^i$ puts strictly positive probability on all pure strategies that yield maximum marginal utility to player i, given she expects to play $\bar{\sigma}^i$. Then, for every σ_ε^{-i} that is "sufficiently close" to σ^{-i}, there exists a mixed strategy PE $\tilde{\sigma}^i$ for player i with $\Gamma(\tilde{\sigma}^i) = \Gamma(\bar{\sigma}^i)$: there exists $\varepsilon > 0$ such that*

$$||\sigma_\varepsilon^{-i} - \sigma^{-i}|| \leq \varepsilon \ \ \Rightarrow \ \ \exists \tilde{\sigma}^i \in R(\sigma_\varepsilon^{-i}) \text{ with } \Gamma(\tilde{\sigma}^i) = \Gamma(\bar{\sigma}^i).$$

First, according to Proposition 3.1(i), if a mixed strategy PE involves player i mixing only over non-redundant pure strategies, no other mixed-strategy PE exists that involves mixing over the same set of pure strategies. The loss-averse player is *decisive*: she has at most one credible mixed plan over a given set of pure strategies. This is in stark contrast to the case of a player with standard expected-utility preferences, who is willing to play any probabilistic mixture over a particular set of pure strategies given that there is at least one incentive compatible mixture. To establish an intuition for this property, recall that for a loss-averse player the expected utility of each pure strategy depends on what she expected to do beforehand. For example, if she thought to play some action with rather high probability, she may become attached to the idea that the associated outcomes will occur and actually prefers to play this action with certainty. Conversely, if a player deemed playing a particular pure strategy rather unlikely, she may favor not to play this action at all. As a consequence, there is a unique intermediate expectation $\hat{\sigma}^i$ such that the player is indeed indifferent between several pure strategies; i.e., the plan $\hat{\sigma}^i$ is the unique mixed strategy over this set of pure strategies that she is willing to follow through.

[14]Proposition 3.1 and the following results show that the implications of loss aversion for strategic interaction and equilibrium play are structurally very different from those of risk aversion under Expected Utility Theory. Since a change in the degree of risk aversion under Expected Utility Theory essentially only changes the entries of the game matrix, the behavioral features that we identify for loss-averse players will never prevail for risk-averse players with standard preferences.

Second, with regard to Proposition 3.1(ii), whenever a mixed strategy PE exists, then for every ε-disturbance of the opponents' strategies another mixed strategy PE over the same set of pure strategies exists. The loss averse-player is *adaptive*, since she is able to adjust expectations as a response to a change in behavior of the opponents, such that she is again willing to play a mixture over the same set of pure strategies. Again, this result is in contrast to the case of a player with standard preferences. Suppose that some mixed strategy is a best-response for player i with standard preferences to a strategy profile σ^{-i} of her opponents. Generically, an arbitrary slight change in the opponents' behavior will alter the marginal material utilities associated with player i's pure strategies. In consequence, after that change player i will not be willing to continue mixing over all those pure strategies that she was willing to play with strictly positive probability before that change. For an expectation-based loss-averse player, however, it is not marginal material utility alone that determines a pure strategy's attractiveness but also marginal psychological utility, where the latter is directly influenced by her expectations regarding her own behavior. According to Proposition 3.1(ii), there exists an adjustment in expectations that exactly offsets the effect of the change in her opponents' behavior on the attractiveness of player i's pure strategies. As a result, she is still willing to mix over the same set of pure strategies.

Illustration: Strategic Behavior for Fixed Expectations

Reconsider the Chicken game introduced at the beginning of this section. According to (3.2), player 1's expected utility of playing strategy $\sigma^1 = (\alpha_1, \alpha_2)$ given the strategy of the opponent, $\sigma^2 = (\beta_1, \beta_2)$, and the fixed expectation about her own strategy, $\hat{\sigma}^1 = (\hat{\alpha}_1, \hat{\alpha}_2)$, is given by

$$\begin{aligned} U^1(\sigma^1, \hat{\sigma}^1, \sigma^2) =& 3\alpha_1\beta_2 + \alpha_2\beta_1 + 2\beta_2\alpha_2 \\ &+\eta\Big[\alpha_1\beta_1[-3\lambda\hat{\alpha}_1\beta_2 - \lambda\hat{\alpha}_2\beta_1 - 2\lambda\hat{\alpha}_2\beta_2] + \alpha_1\beta_2[3\hat{\alpha}_1\beta_1 + 2\hat{\alpha}_2\beta_1 + \hat{\alpha}_2\beta_2] \\ &\quad + \alpha_2\beta_1[\hat{\alpha}_1\beta_1 - 2\lambda\hat{\alpha}_1\beta_2 - \lambda\hat{\alpha}_2\beta_2] + \alpha_2\beta_2[2\hat{\alpha}_1\beta_1 - \lambda\hat{\alpha}_1\beta_2 + \hat{\alpha}_2\beta_1]\Big], \end{aligned}$$

where the first line describes expected material utility, whereas the remaining lines capture gain-loss utility. The latter contains comparisons of each outcome of the actual lottery to every outcome of the reference lottery, weighted with the respective occurrence probabilities. Note that the marginal expected utilities

$$\frac{\partial U^1(\sigma^1, \hat{\sigma}^1, \sigma^2)}{\partial \alpha_1} = 3\beta_2 + \eta\left\{\beta_2^2\hat{\alpha}_2 + \beta_1\beta_2(1-\lambda)[3\hat{\alpha}_1 + 2\hat{\alpha}_2] - \lambda\hat{\alpha}_2\beta_1^2\right\} \tag{3.7}$$

$$\frac{\partial U^1(\sigma^1, \hat{\sigma}^1, \sigma^2)}{\partial \alpha_2} = 1 + \beta_2 + \eta\left\{\hat{\alpha}_1\beta_1^2 + \beta_1\beta_2[(1-\lambda)\hat{\alpha}_2 + 2\hat{\alpha}_1] + \beta_2^2(1-\lambda)\hat{\alpha}_1\right\} \tag{3.8}$$

do not depend on the probabilities α_1 and α_2 of player 1 actually going straight or actually swerving, respectively. Hence, if the marginal utility in (3.7) is larger than the one in (3.8)

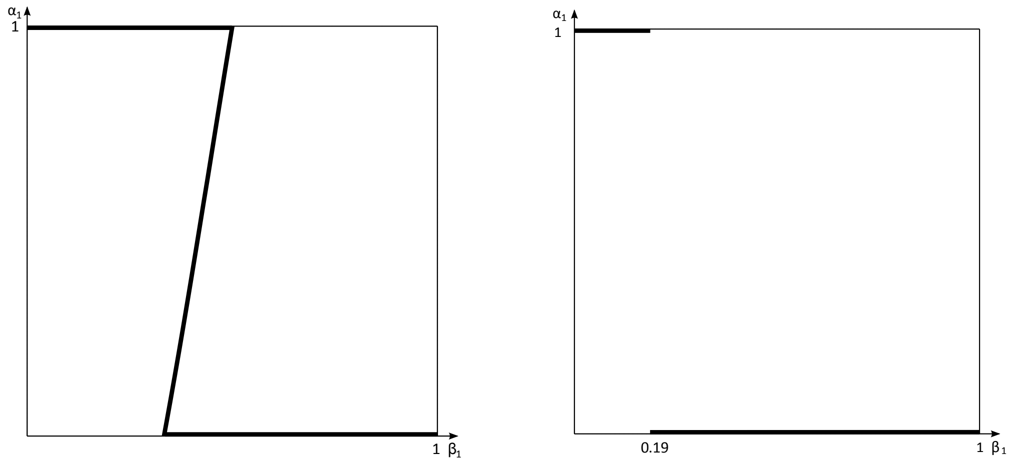

Figure 3.2.: Player 1's set of PEs and set of CPEs if $\lambda = 3$ and $\eta = 1$.

given that player 1 expects to go straight for sure ($\hat{\alpha}_1 = 1$), actually going straight for sure ($\alpha_1 = 1$) is a PE. Here, player 1 seeks to increase α_1 as much as possible and she is willing to follow through the plan of going straight for sure. Likewise, swerving for sure ($\alpha_2 = 1$) is a PE if the marginal utility in (3.7) is smaller than the one in (3.8) given that player 1 expects to swerve for sure ($\hat{\alpha}_2 = 1$). Finally, for a mixed strategy with $0 < \alpha_1, \alpha_2 < 1$ to be a PE, the marginal utilities in (3.7) and (3.8) have to be identical for expectations $\hat{\alpha}_1 = \alpha_1$ and $\hat{\alpha}_2 = \alpha_2$.

The left panel of Figure 3.2 depicts (for $\lambda = 3$ and $\eta = 1$) the set of values of α_1 that constitute a PE for player 1 in response to player 2 going straight with probability β_1. Similar to a player with standard expected-utility preferences, an expectation-based loss-averse player has a unique PE if her opponent plays a particular pure strategy with rather high probability. Given player 2 almost surely swerves (goes straight), the only expectation that player 1 indeed follows through is to go straight (swerve).

However, the two qualitative differences from Proposition 3.1 also become apparent in Figure 3.1. First, while there exists a unique value of β_1 such that a player with standard preferences is indifferent between all probabilistic mixtures over the two pure strategies, the expectation-based loss-averse player is *decisive*: there is at most one mixed strategy that she may play in response to a particular value of β_1. If player 1 thought to swerve with a high probability, she becomes attached to the idea that no crash will occur and actually prefers to swerve with certainty. Conversely, if she thought that she will most likely go straight, she relishes the idea of becoming a local hero and indeed favors to go straight with certainty. To comprehend this attachment effect, note that $\frac{d^2U^1(\sigma^1,\hat{\sigma}^1,\sigma^2)}{d\alpha_j d\hat{\alpha}_j} \geq 0$, $j \in \{1, 2\}$. Hence, the attractiveness of a pure strategy is increasing in the expectation to play this strategy. As a consequence, there is a unique intermediate expectation $\hat{\alpha}_1$ such that she is indeed indifferent between both pure strategies and this plan alone constitutes a mixed PE.

Second, if player 1 has standard expected-utility preferences, she is willing to mix only if player 2 goes straight with probability $\beta_1 = 1/2$. If player 1 is expectation-based loss-averse, however, she is *adaptive*. Hence, there exists a non-trivial range $[\underline{\beta}, \frac{1}{2}]$ such that player 1 may play some mixed strategy in response to any $\beta_1 \in [\underline{\beta}, \frac{1}{2}]$. Consider a mixed strategy with $\beta_1 \in \left(\underline{\beta}, \frac{1}{2}\right]$ for player 2 such that for player 1 exactly one mixed PE, denoted by $\tilde{\sigma}^1$, exists. Suppose player 2 slightly reduces her probability to go straight, i.e., β_1 decreases. Going straight then becomes more attractive for player 1 as it is associated with a higher probability to be the public hero. In consequence, $\tilde{\sigma}^1$ no longer constitutes a credible plan because deviating by going straight with certainty is profitable. However, expecting to swerve with a higher probability re-attaches player 1 to swerving and makes both pure strategies equally attractive again. Hence, an adjusted credible mixed plan with a lower probability to go straight exists and player 1 remains willing to mix over the same set of pure strategies.

3.5.3. Choice-Acclimating Personal Equilibrium

For this section we assume that each player's expectation regarding her own behavior is not fixed when she takes her action but pinned down by the action taken. As a consequence, the lottery over material utility outcomes induced by player i's actual action coincides with the reference lottery over material utility outcomes that player i expected. In this case, it turns out that a player always prefers not to play a mixed strategy. Consider two distinct (possibly mixed) strategies of player i that induce different lotteries over material outcomes. By Lemma 3.1, mixing between these two strategies creates an additional degree of riskiness, implying a negative effect on psychological utility. Hence, a player always prefers to play one of the two strategies with certainty over mixing between them.

In order to understand the most basic driving forces of player i's strategic behavior in this case, consider the following situation: There is no move of nature and all of player i's opponents play pure strategies. Player i's material utility outcome from playing a particular pure strategy is therefore deterministic. If player i randomizes between two pure strategies which result in different utility outcomes, the comparison of the material utility outcomes results in a net loss. Now, consider a deviation from this mixed strategy to one of the pure strategies. As the player receives exactly the material utility outcome she expected to obtain, this eliminates any net losses, thereby making the mixture over the pure strategies rather unattractive.

More generally, reducing the number of pure strategies that player i mixes over favorably affects the gain-loss utility by reducing the number of outcome comparisons and thus the number of net losses that reduce expected utility. This intuition is formally reflected in the following proposition, which documents a general *reluctance to mix* in CPE situations.

Proposition 3.2. ***Reluctance to mix:*** *Suppose $\bar{\sigma}^i$ with $|\Gamma(\bar{\sigma}^i)| \geq 2$ is a CPE for player i given* *the strategy profile σ^{-i} of her opponents. Then*

$$L^i(s', \sigma^{-i}) = L^i(s'', \sigma^{-i}) \quad \forall s', s'' \in \Gamma(\bar{\sigma}^i).$$

A loss-averse player is willing to mix only between pure strategies that induce identical lotteries over material utility outcomes. Since mixing over such strategies results in a lottery that is not different from the lottery over material utility outcomes induced by the pure strategies, the player is willing to play a mixed strategy only if mixing has no effect.

Illustration: Strategic Behavior for Choice-Acclimating Expectations

The expected utility of playing (and expecting to play) σ^1 given σ^2 is

$$\begin{aligned} U^1(\sigma^1, \sigma^1, \sigma^2) = {} & 3\alpha_1\beta_2 + \alpha_2\beta_1 + 2\alpha_2\beta_2 \\ & - \eta(\lambda - 1)\Big[3\alpha_1^2\beta_1\beta_2 + 2(2\alpha_1\alpha_2\beta_1\beta_2) + \alpha_1\alpha_2\beta_1^2 + \alpha_1\alpha_2\beta_2^2 + \alpha_1^2\beta_1\beta_2\Big]. \end{aligned}$$

Given player 2 goes straight with probability β_1, going straight (and expecting to go straight) with probability α_1 is a CPE for player 1 if this maximizes expected utility $U^1(\sigma^1, \sigma^1, \sigma^2)$. The set of CPEs is depicted in the right panel of Figure 3.2. As for a player with standard expected-utility preferences, swerving (going straight) for sure is the unique CPE given the other player rather likely goes straight (swerves). Unlike a player with standard expected-utility preferences, however, she never deliberately plays a mixed strategy. Even if player 1 is indifferent between playing (and expecting to play) either one of the two pure strategies, she incurs a strictly lower expected utility from any mixture of these. Playing a mixture creates "additional" uncertainty about material utility outcomes and, thus, net losses. Note that for $0.19 < \beta_1 < 0.5$ the loss-averse player prefers to swerve for sure although the expected material utility favors going straight. To understand this, note that the average self-distance of the lottery induced by going straight strictly exceeds the one induced by swerving. This implies a lower psychological utility from going straight compared to swerving, which needs to be outweighed by a higher expected material utility to make the loss-averse player willing to go straight.

3.6. Equilibrium Existence and Behavior

Section 3.5 demonstrated how the strategic behavior of loss-averse players differs from the behavior of their counterparts with expected-utility preferences. In the following, we discuss the resulting implications for equilibrium behavior and equilibrium existence for the notions of PNE and CPNE as introduced in Section 3.4.

We start with the simplest case in which the game is free of any inherent uncertainty, i.e., $\Theta = \{\tilde{\theta}\}$. In this setting, the set of pure strategy Nash equilibria is identical to the set of pure strategy PNEs and also the set of pure strategy CPNEs. Consider PNE first. Given a player expects to play the pure strategy Nash best response to a given pure strategy profile of her opponents, any deviation results in not only (weakly) lower material utility but in addition creates unexpected losses—and therefore is not profitable. Conversely, expecting to play a pure strategy that is not a Nash best response cannot constitute a PE, because the deviation to the Nash best response would yield not only a strictly higher deterministic material utility payoff but also—due to the unexpected gain—strictly higher psychological utility. Hence, for a given pure strategy profile of her opponents, a particular pure strategy is a PE for player i if and only if it is a Nash best response. Since these considerations apply to each player, the identity of the set of Nash equilibria and the set of PNEs follows immediately. For choice-acclimating expectations the case is even more apparent. Since no uncertainty is involved in the game as long as the players play pure strategies, there are no gains or losses involved for a player whose expectations match actual behavior. Hence, her utility from playing any pure strategy is identical to the utility of a player with standard preferences. Together with the reluctance to deliberately play mixed strategies (cf. Proposition 3.2), it follows that the set of pure strategy CPNEs is also identical to the set of pure strategy Nash equilibria.

Proposition 3.3. *Suppose there is no inherent uncertainty in the game, $\Theta = \{\tilde{\theta}\}$. Then the following statements are equivalent:*

(i) $s \in \mathcal{S}$ *is a Nash equilibrium.*

(ii) $s \in \mathcal{S}$ *is a CPNE.*

(iii) $s \in \mathcal{S}$ *is a PNE.*

We conclude that in simple games without uncertainty—e.g, the Chicken game, the Prisoners Dilemma, or the Battle of the Sexes—it is possible that loss-averse players behave *as if* they had standard preferences. With regard to pure strategy equilibria, the equilibrium behavior of expectation-based loss-averse players even is necessarily identical to the behavior in Nash equilibria. This picture, however, changes if there is either uncertainty in the game or if mixed strategies are taken into account.

3.6.1. Personal Nash Equilibrium

As we have seen in Section 3.5, the existence of two or more pure strategy PEs for a given strategy profile of the opponents does not imply that every mixture over these pure strategies is also a PE. Instead, decisiveness implies that there exists at most one such mixture constituting a mixed strategy PE—cf. Proposition 3.1(i). Therefore, the PE correspondences are not convex valued, Kakutani's fixed point theorem is not applicable, and the existence of PNEs is a priori unclear.

Nevertheless, we can establish the existence of a PNE and pin down equilibrium play for two basic cases. First, if there exists a Nash equilibrium in (materially) weakly dominant pure strategies, this constitutes also the unique PNE. This finding is rooted in the fact that it is always a credible plan to expect to play a (materially) weakly dominant pure strategy. Here, strategy s^i is (materially) weakly dominant if $u^i((s^i, s^{-i}), \theta) \geq u^i((\tilde{s}^i, s^{-i}), \theta)$ for all pure strategy profiles $(\tilde{s}^i, s^{-i})$ and all states of the world θ, where for each $(\tilde{s}^i, s^{-i})$ the inequality is strict for at least one θ.[15] Intuitively, deviating to a dominated strategy $\tilde{s}^i$ not only reduces expected material utility, but, given that the reference lottery over outcomes is induced by the dominant strategy s^i, also reduces gains (or turns them into losses) and increases losses.

Second, for games with two players each of whom has two actions the existence of a PNE is guaranteed. If for a given strategy of her opponent each of a player's two pure strategies constitutes a PE, there also exists a mixed strategy PE. Essentially, when the strategy of the opponent changes, adaptiveness induces this mixed strategy PE to change continuously thereby providing a connection between the sets of pure strategy PEs. Thus, a player's PE correspondence has a connected graph. Furthermore, this PE correspondence has full support over the strategy space of the player's opponent.[16] In consequence, a PNE must exist.

Proposition 3.4. *Regarding PNE, the following statements hold:*

(i) Suppose (s^i, s^{-i}) is a Nash equilibrium in (materially) weakly dominant strategies. Then (s^i, s^{-i}) is the unique PNE.

(ii) Suppose $\mathcal{I} = \{1, 2\}$ and $|\mathcal{S}^i| = 2$ for $i = 1, 2$. Then there exists a PNE.

Proposition 3.4(i) derives the PNE for several prominently studied games. For example public good games with monetary, and thus discrete, contributions and payoffs always have a PNE—even if there is uncertainty about the other players' endowment. More specifically, the tendency to free ride and not to contribute remains an equilibrium also under loss aversion. Similarly, in the Vickrey auction with monetary bids and valuations it is a PNE to bid the true valuation for loss-averse players.

Illustration: Equilibrium Behavior for Fixed Expectations

Reconsider the Chicken game introduced in Section 3.5. The middle panel of Figure 3.3 depicts the sets of PEs for both players. According to Definition 3.1, the game's PNEs lie at the intersections of the two PE correspondences. As implied by Proposition 3.3, the Nash equilibria in

[15]This definition of dominance is based upon the idea that nature can be interpreted as an additional player in the game. Hence, if a strategy is weakly dominant for player i, it provides weakly higher utility than any other of her feasible strategies irrespectively of the opponents' strategies and nature's draw.

[16]See Theorem 1 (p. 422) in Kőszegi (2010).

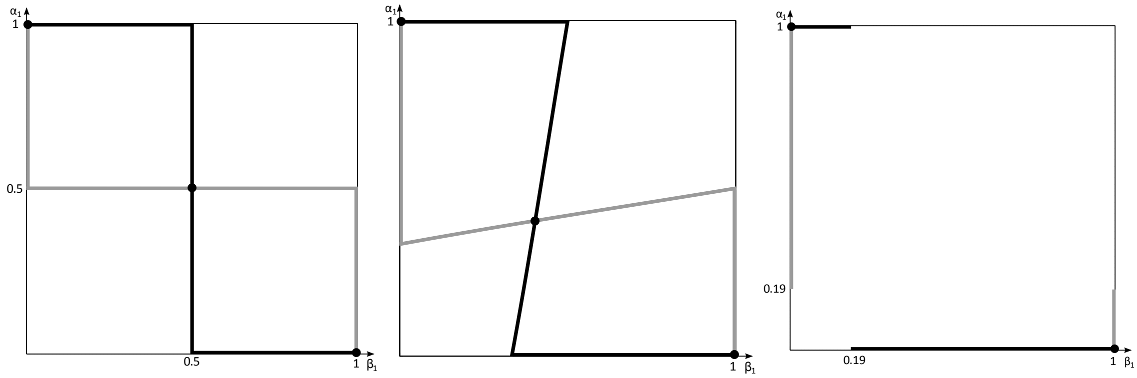

Figure 3.3.: Nash equilibria for players with standard preferences, PNEs and sets of CPNEs if $\lambda = 3$ and $\eta = 1$.

pure strategies also constitute PNEs. Thus, the game has two pure strategy PNEs in each of which one driver goes straight for sure and the other driver swerves for sure. Furthermore, there also exists a mixed-strategy PNE which has both drivers going straight with a 40% chance and swerving with a 60% chance. Obviously, this mixed strategy PNE differs from the mixed strategy Nash equilibrium, which has both players going straight with a 50% chance. If $\beta_1 = 0.5$, expected material utility of both actions is identical for player 1. In this case, the option to go straight is more risky, though. In particular, the lottery over material utility outcomes induced by going straight for sure is a mean preserving spread of the one that is induced by swerving for sure. Since a loss-averse player tends to avoid risks, player 1 is not willing to go straight with positive probability if $\beta_1 = 0.5$ but only if player 2 is sufficiently more likely to swerve than to go straight, which increases the expected material utility from going straight over the expected material utility from swerving. Overall, this implies a PNE in which the more "risky" option is associated with higher expected material utility.

3.6.2. Choice-Acclimating Personal Nash Equilibrium

In Proposition 3.2, we identified a general reluctance of agents with choice-acclimating expectations to deliberately randomize between pure strategies with different probabilistic consequences. This behavioral feature immediately implies that a mixed strategy CPNE can only exist if for some player two of her pure strategies lead to identical probabilistic consequences.

Corollary 3.1. *Suppose that for any player $i \in \mathcal{I}$ and any strategy profile σ^{-i} of her opponents each two pure strategies induce different lotteries over material utility outcomes, i.e., $L^i(s^i_k, \sigma^{-i}) \neq L^i(s^i_m, \sigma^{-i})$ for all $i \in \mathcal{I}$, $\sigma^{-i} \in \Sigma^{-i}$, and $s^i_k, s^i_m \in \mathcal{S}^i$ with $s^i_k \neq s^i_m$. Then the following statements hold.*

(i) A mixed strategy CPNE does not exist.

(ii) For $\Theta = \{\tilde{\theta}\}$, if there exists no pure strategy Nash equilibrium, there exists no CPNE.

Corollary 3.1(i) complements papers that restrict attention to pure strategy CPNEs when studying the strategic interaction of expectation-base loss-averse players by showing that the focus on pure strategy equilibria is without loss of generality. The result can be applied to a large variety of settings. For example, there is no CPNE in which agents randomly choose their efforts in the team production setting of Daido and Murooka (2014) or in any finite version of the rank-order tournaments studied in Gill and Stone (2010) and Dato, Grunewald, and Müller (2015a). Likewise, bidders never deliberately randomize over their bids in any finite version of the auctions analyzed in Lange and Ratan (2010) and Eisenhuth (2010).[17] Corollary 3.1(ii), which follows from Propositions 3.2 and 3.3, also implies that a CPNE does not exist in all settings. Take for example a slightly asymmetric Matching Pennies game that has no redundant strategies, no inherent uncertainty, and no pure strategy Nash equilibrium. Even for this basic game a CPNE does not exist because loss-averse players do not deliberately mix over pure strategies. This strongly suggests that with regard to CPNE the question of existence has to be investigated in any application. According to Corollary 3.1(i), however, this investigation can be restricted to the question of the existence of pure strategy CPNEs.

While existence of a CPNE is not guaranteed, we can establish sufficient conditions for a CPNE to exist and identify equilibrium play in these cases. First, Proposition 3.3 yields a very simple sufficient condition for the existence in games without inherent uncertainty: if a pure strategy Nash equilibrium exists, it is also a CPNE. Second, equilibrium play in and existence of a CPNE can also be linked to the existence of a Nash equilibrium in (materially) weakly dominant pure strategies. Unlike to the case where expectations are fixed, playing a (materially) weakly dominant pure strategy not necessarily constitutes a CPE. The reason is that, given a player plays some (materially) weakly dominated strategy, the reference lottery is also induced by this strategy, which in fact may lead to a smaller net loss than the (materially) weakly dominant strategy. However, as long as the weight that the player attaches to this net loss does not exceed the weight on material utility, i.e., $\eta(\lambda - 1) \leq 1$, the higher expected material utility associated with the (materially) weakly dominant strategy outweighs any potential reduction in psychological utility.

Proposition 3.5. *Suppose that $\eta(\lambda - 1) \leq 1$ and that (s^i, s^{-i}) is a Nash equilibrium in (materially) weakly dominant strategies. Then (s^i, s^{-i}) is the unique CPNE.*

[17] To be precise, most of the above applications comprise multidimensional outcomes. In Proposition 3.6, we show that our results carry over to the case of multidimensional outcomes.

Illustration: Equilibrium Behavior for Choice Acclimating Beliefs

The right panel of Figure 3.3 shows the set of CPEs of the two drivers in the Chicken game. As implied by Proposition 3.3, the two pure strategy Nash equilibria of the game also constitute CPNEs. Moreover, Figure 3.3 also illustrates the non-existence of a mixed strategy CPNE, which is rooted in a loss-averse player's reluctance to deliberately mix over pure strategies in CPE situations. Even if her opponent plays a mixture between swerving and going straight that induces both actions to be a CPE for a player, she would not be willing to mix between these two pure strategies. Therefore, in contrast to the case with standard preferences or to situations with fixed expectations, there exist only two CPNEs in the Chicken game.

3.7. Discussion

3.7.1. Interpretation of Mixed Strategies and Equilibrium Existence

So far, we have seen that the existence of PNE is a priori not clear and the existence of CPNE may fail even in simple games. Importantly, the possible nonexistence of equilibria relies on the notion that each individual player indeed mixes over her pure strategies. In the last decades, however, there have emerged different views on how to interpret mixed strategies. For example Aumann and Brandenburger (1995) argue that even if every player chooses a definite action other players may not know which one. In their interpretation a probabilistic mixture represents a players' conjecture about her opponents' choices and not randomness in her opponents' strategies. Adopting this notion, a CPNE and a PNE in conjectures necessarily exists. To see this, suppose a player is indifferent between several pure strategies. As her opponents do not know which of these she will play, their conjectures can involve each of these pure strategies and all mixtures between them. As a consequence, the set of feasible conjectures is the convex hull of the set of best responses. Due to the continuous differentiability of utility functions, Kakutani's fixed point theorem then is applicable and an equilibrium exists.

Along similar lines, Rosenthal (1979) proposes to interpret players not as individuals per se but as large populations of individuals. In a game, randomly drawn individuals, one from each such population, play against each other. In the large population represented by player i a mixture over pure strategies thus is not necessarily generated by individual mixing but may also reflect the distribution of pure strategy choices in that population. If the distributions over pure strategy choices in the populations represented by player i's opponents induce the existence of several pure strategy best replies for the individuals in the population represented by player i, each of those individuals is willing to play either one of these best replies. Therefore, when playing against a random draw from player i's population, the individuals in her opponents' populations can in turn rationally expect to face any mixture between the respective pure strategy best replies. In consequence, playing against a large population is as if playing against a single player that is additionally willing to play any mixture between pure strategy best replies. We conclude that there exists a "large population" equilibrium if the convex hulls of the best response correspondences intersect. This is again guaranteed by Kakutani's fixed point theorem such that both a CPNE and a PNE always exist when the large-population interpretation of mixed strategies is applied.[18]

Remark 1. *Following the reinterpretation of mixed strategies proposed by Aumann and Brandenburger (1995) a PNE and a CPNE in conjectures always exist. Similarly, a large population PNE and CPNE à la Rosenthal (1979) always exist.*

3.7.2. Multidimensional Outcomes

Often material outcomes comprise multiple consumption dimensions. For example, winning an auction may come along with a gain in the good dimension from obtaining the object that was for sale and a loss in the money dimension from having to pay the winning bid. Therefore, an important aspect of the behavior of loss-averse agents is how they deal with multidimensional outcomes, in which case a single outcome may simultaneously generate gains and losses along different dimensions. In this section, we show that our results carry over to the case of multidimensional outcomes. Each player $i \in \mathcal{I}$ has payoff function $u^i : \mathcal{S} \times \Theta \to \mathcal{U}^i \subset \mathbb{R}^R$ which maps any combination of a pure strategy profile $s \in \mathcal{S}$ and a random realization of θ into a payoff vector which comprises $R \geq 2$ different consumption dimensions, $u^i(s,\theta) = (u^i_1(s,\theta), \ldots, u^i_R(s,\theta)) \in \mathbb{R}^R$. $P^i(u|\sigma)$ then describes the probability that utility vector u is realized for player i under the strategy profile σ. Following Kőszegi and Rabin (2006), material utility and gain-loss utility are assumed to be additively separable over dimensions, yielding overall utility

$$\begin{aligned} U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i}) = & \sum_{u \in \mathcal{U}^i} P^i(u|(\sigma^i, \sigma^{-i})) \cdot \sum_{r=1}^{R} u_r \\ & + \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|(\sigma^i, \sigma^{-i})) \cdot P^i(\tilde{u}|(\hat{\sigma}^i, \sigma^{-i})) \cdot \sum_{r=1}^{R} \mu(u_r - \tilde{u}_r). \end{aligned} \tag{3.9}$$

[18] For an example how these interpretations generate additional equilibria reconsider the large population interpretation for choice acclimating beliefs in the Chicken game. When applying this logic, there is a third equilibrium which lies at the intersection of the convex hulls of the sets of CPEs. For every individual in the population corresponding to player 1 going straight and swerving are pure strategy CPEs if 19% of the individuals in the population representing player 2 go straight and 81% swerve. In this case, every single individual in player 1's population may either swerve for sure or go straight for sure. This, in turn, implies that population shares for player 1 that go straight and swerve, respectively, can be exactly such that for each individual in player 2's population swerving for sure and going straight for sure are both pure strategy CPEs. Overall, this leads to a "large population" CPNE $(\alpha_1, \alpha_2) = (\beta_1, \beta_2) = (0.19, 0.81)$.

The gain-loss utility from multidimensional outcome u when having expected $\tilde{u}$ is determined by comparing material utilities for each dimension separately.[19] Thus, a particular outcome may give rise to mixed feelings if it is associated with losses in some dimensions and with gains in other dimensions.

Nevertheless, the definition of a *redundant* pure strategy directly carries over to the case of multidimensional outcomes. Moreover, in case of multidimensional outcomes, we define a pure strategy to be (weakly) materially dominant if the strategy is (weakly) materially dominant in every dimension. With these slightly amended definitions, the results from Sections 3.5 and 3.6 carry over to the case of multidimensional payoffs.

Proposition 3.6. *Suppose material payoffs are multidimensional. Then the results from Proposition 3.1, Proposition 3.2, Proposition 3.4, Corollary 3.1 and Proposition 3.5 continue to hold.*

The fact that our results also hold for multidimensional outcomes is rooted in the separability of utility across dimensions. Adding payoff dimensions does not eliminate but rather strengthens the effects of loss aversion. With regard to the basic case of games without inherent uncertainty, this implies that players with fixed expectations get attached even more strongly to their plans. Consequently, as the following generalization of Proposition 3.3 shows, more outcomes can be supported in equilibrium.

Proposition 3.7. *Suppose there is no draw of nature, $\Theta = \{\tilde{\theta}\}$, and all players' payoffs are multidimensional, $\mathcal{U}^i \subset \mathbb{R}^R$ with $R \geq 2$ for all $i \in \mathcal{I}$. Then the following statements hold:*

(i) $s \in \mathcal{S}$ *is a CPNE if and only if it is a NE.*

(ii) $s \in \mathcal{S}$ *is a PNE if it is a NE.*

(iii) A pure-strategy profile $s \in \mathcal{S}$ is implementable as PNE for λ sufficiently large if for each $\tilde{s}^i \neq s^i$, $i \in \mathcal{I}$, there exists some dimension $r^i(\tilde{s}^i) = 1, \ldots, R$ such that $u^i_{r^i(\tilde{s}^i)}((s^i, s^{-i}), \tilde{\theta}) > u^i_{r^i(\tilde{s}^i)}((\tilde{s}^i, s^{-i}), \tilde{\theta})$.

Without inherent uncertainty in the game, the logic underlying parts (i) and (ii) of Proposition 3.7 is the same as for the corresponding statements regarding one-dimensional payoffs in Proposition 3.3. In contrast to the case of one-dimensional payoffs, however, under multidimensional payoffs a pure strategy combination that is not a Nash equilibrium might form a PNE—cf. Proposition 3.7(iii). A deviation from some pure strategy yielding lower material utility in at least one dimension creates a loss and thus, it is unattractive for a sufficiently strong degree of loss aversion even if it increases overall material utility. As a consequence, every pure strategy combination such that for every player $i \in \mathcal{I}$ any unilateral deviation yields lower material utility in at least one consumption dimension can be supported in a PNE. This reveals

[19]Here we assume a universal gain-loss function $\mu(\cdot)$ that applies to all consumption dimensions. Allowing for dimension-specific gain-loss functions $\mu_1(\cdot), \ldots, \mu_R(\cdot)$ would not change our results qualitatively.

that the common practice in standard game theory to consolidate different consumption dimensions is not without loss of generality if players are loss averse because PNEs are potentially eliminated.

3.8. Conclusion

This paper provides a comprehensive analysis of the strategic interaction of expectation-based loss-averse players. Taking mixed strategies into account, we show how the equilibrium concepts of Kőszegi and Rabin (2006, 2007) are applicable to strategic multi-player settings. For loss-averse players the attractiveness of pure strategies is directly influenced by their expectations and, thus, a player's expected utility is not linear in the mixing probabilities she assigns to her pure strategies. Expectation-based loss-averse players differ in their strategic behavior from players with standard expected-utility preferences in several respects. First, for fixed expectations, loss-averse players are *adaptive* in the sense that mixed strategies may be part of a "best" response of a player for a nontrivial range of opponents' strategies. Second, loss-averse players are *decisive* with respect to mixed strategies, i.e., for given strategies of the opponents there is at most one mixed "best response". Third, for choice-acclimating expectations, loss-averse players are *reluctant* to play mixed strategies irrespective of the game.

The strategic behavior has direct implications for resulting equilibria. In two basic cases loss aversion does not affect equilibrium play compared to standard expected utility: first, if there is no inherent uncertainty in the game under consideration and payoffs are one-dimensional; second, if the game is solvable in weakly dominant strategies. This picture changes as soon as either mixed strategy equilibria are studied or uncertainty is involved. In particular, if expectations are choice acclimating, mixed strategy equilibria never exist. If expectations are fixed, on the other hand, players get attached to the strategy that they expected to play even if randomness is involved in the strategy. Thus, mixed strategy equilibria may exist in this case.

This paper paves the way to a variety of further research questions. First, we showed that loss aversion may increase the number of equilibria, particularly if payoffs are multidimensional. Extending the selection criterion *preferred personal equilibrium (PPE)*, as proposed in Kőszegi and Rabin (2006), to strategic interaction, however, is not as promising as a cursory first glance seems to suggest. More specifically, while it is evident that at least one PE provides maximal expected utility for the individual decision context, it may well be the case that there does not exist a combination of strategies such that all players play their most preferred PE given the other players' strategies. Thus, it may well be that no PNE survives the straightforward application of PPE to strategic interaction.[20] It seems interesting—if not necessary—to investigate sensible criteria for equilibrium selection for the equilibrium concepts proposed in Section 3.4.

[20]As an example consider a slightly asymmetric matching pennies game, in which the only PNE involves mixed

Second, we study games with finite action spaces. Some interesting applications like auctions or tournaments, however, involve continuous choice variables like effort choices or money bids, respectively. Although the intuition behind the resulting strategic behavior should be similar in spirit to the insights gathered in this paper, the technical apparatus involved in the derivation is somewhat different. An extension of our results regarding mixed strategies would allow a more comprehensive study of equilibria in these contexts.

strategies. As shown in the proof of Proposition 3.4, given the opponent plays her equilibrium strategy, the two pure strategies are also PEs. Since at least one of their pure strategies must constitute a CPE, the mixed PE can never be a PPE and the PNE is not a mutual preferred personal Nash equilibrium.

4. Expectation-Based Loss Aversion and Rank-Order Tournaments

4.1. Introduction

Relative performance evaluation in the form of rank-order tournaments is commonplace—electoral competition in politics, contests in professional sport, or promotion tournaments within a particular corporation or the labor market in general. In the light of the widespread applicability of rank-order tournaments, it is hardly surprising that, beginning with the seminal article by Lazear and Rosen (1981), economic scholars have studied the strategic interaction of the contestants participating in this particular form of incentive mechanism for over three decades.

Recently, several theoretical contributions have enriched the canonic tournament model by incorporating insights gathered in the psychological or experimental economics literature.[1] Gill and Stone (2010) assume that the preferences of contestants are reference dependent and exhibit loss aversion.[2] Applying the concept of choice-acclimating personal equilibrium (CPE) as defined in Kőszegi and Rabin (2007), they assume that a contestant's reference point is determined by his rational expectations about outcomes, where these expectations correctly incorporate the effect of his behavior.[3,4] Interestingly, Gill and Stone (2010) find that if the degree of loss

[1] For example, Santos-Pinto (2010) as well as Ludwig, Wichardt, and Wickhorst (2011) investigate overconfidence of contestants. Furthermore, contestants' preferences have been modified to capture inequity aversion (Grund and Sliwka, 2005b; Demougin and Fluet, 2003) or joy of winning (Kräkel, 2008).

[2] Reference-dependent preferences and loss aversion have been introduced into the economic discourse by Kahneman and Tversky (1979) as key aspects of their prospect theory.

[3] More specifically, following the disappointment aversion concept in Bell (1985), Gill and Stone (2010) posit that a contestant's reference point in the money dimension corresponds to the average prize that he will receive in the tournament given his own and his opponent's effort choice. Since the tournament outcome is binary—i.e., a contestant either wins or loses—this formulation is equivalent to the notion of CPE as introduced in Kőszegi and Rabin (2007).

[4] Evidence supporting the reference point formation according to Kőszegi and Rabin is provided by Abeler, Falk, Goette, and Huffman (2011), Crawford and Meng (2011), Ericson and Fuster (2011), and Gill and Prowse

aversion among homogeneous contestants is strong, the symmetric equilibrium ceases to exist. Intuitively, uncertainty about the tournament outcome is maximized in a symmetric equilibrium in the sense that each contestant faces a 50% probability of being victorious. As expectation-based loss-averse contestants strongly dislike this uncertainty, in equilibrium they exert rather different levels of effort, which lead to unequal winning probabilities and thereby reduce uncertainty. Overall, with reference dependence and loss aversion being widely recognized as relevant determinants of individual risk preferences, this result is an important caveat to many of our insights from tournament theory, which typically rest upon the existence of symmetric equilibria.

In this note, we show under which conditions a focus on symmetric equilibria in rank-order tournaments with homogeneous contestants seems justifiable even if contestants are expectation-based loss averse. To this end, we provide two arguments that bolster the existence of symmetric equilibria. First, for contestants with choice-acclimating expectations as in Gill and Stone (2010), we argue that a symmetric equilibrium exists for moderate degrees of loss aversion. In fact, for a symmetric equilibrium not to exist, as long as a minimal amount of effort comes without costs, the contestants' concern for psychological gain-loss utility must outweigh their concern for consumption utility. Such a strong degree of loss aversion, however, would also imply violations of stochastic dominance (Kőszegi and Rabin, 2007).

Second, we reconsider behavior in rank-order tournaments if contestants' reference points are shaped by their lagged—i.e., choice-unacclimating—expectations. Under this alternative assumption, there exist symmetric equilibria *for all* degrees of loss aversion. Formally, instead of CPE, we apply the concept of personal equilibrium (PE), as proposed in Kőszegi and Rabin (2006). Hence, the contestants enter the tournament with a given "strategy" or "game plan" in mind, such that action is taken for a fixed set of expectations. Contestants having fixed expectations seems particularly plausible in multistage tournaments, where contestants may ponder their overall chances of winning when entering the tournament. In this case, they will surely enough enter every stage, except for maybe the first one, with fixed expectations. Even in a one-shot tournament, however, fixed expectations may prevail as is suggested by coaches having to announce their teams' rosters to the press hours before the start of tonight's game or candidates for political office selecting their staff long before the actual campaign.

For fixed expectations, there always exist multiple combinations of individual plans that contestants are willing to follow through—both symmetric and asymmetric ones. In order to investigate the robustness of symmetric equilibria for fixed expectations, we address the question of equilibrium selection by employing the equilibrium refinement of preferred personal equilibrium (PPE), as proposed in Kőszegi and Rabin (2006). Similar to the case of CPE, only strong loss aversion may eliminate the symmetric equilibrium under this equilibrium notion. In par-

(2012).

ticular, the degree of loss aversion necessary for the symmetric equilibrium ceasing to exist is even stronger under PPE than under CPE. Our findings thus resonate well with the observation in Kőszegi and Rabin (2007) that CPE embodies a stronger notion of risk aversion than PE: While the strong dislike of uncertainty under choice-acclimating beliefs leads to an asymmetric equilibrium with a fairly certain tournament outcome, the lagged expectation to compete in a balanced tournament with a rather uncertain outcome can in fact be a credible plan.

In order to present these findings as concise as possible, we make use of a streamlined version of the canonic tournament model. More specifically, in our model agents' effort choices affect the probability distribution over output levels but not the output level itself. We thus use a parameterized distribution formulation to set up our tournament environment rather than a state-space formulation as done by Lazear and Rosen (1981). Next to tractability, this modeling choice—which according to Hart and Holmström (1987) often "yields more economic insights" (p.78)—is primarily an educational one, as it allows to delineate how expectation-based loss aversion enriches the strategic interaction of contestants.[5] Nevertheless, to guarantee a comparison of results on a level playing field, we fully replicate the findings by Gill and Stone (2010) with regard to tournaments with homogeneous contestants.

The rest of the paper is organized as follows. In Section 4.2, we introduce the tournament environment and contestants' reference-dependent preferences. This model is analyzed for choice-acclimating expectations in Section 4.3 and for lagged fixed expectations in Section 4.4. After addressing equilibrium selection in Section 4.5, we conclude in Section 4.6.

4.2. The Model

Two agents A and B compete in a rank-order tournament with winner prize W and loser prize $L < W$. Both agents share the Bernoulli utility function $u(x)$ for money, where $u'(x) > 0$. Let $\Delta = u(W) - u(L)$. Agent i can exert effort $e_i \in [0, 1]$ at cost $c(e_i)$, where $c(0) = c'(0) = 0$, $c'(e) > 0$ for $e > 0$, $c''(e) > 0$, $c'''(e) > 0$, and $\lim_{e\to 1} c'(e) = \infty$.[6] Given agent i exerts effort e_i, she produces high output $\pi_i = \overline{\pi}$ with probability e_i and low output $\pi_i = \underline{\pi}$ with probability $1 - e_i$, where $\overline{\pi} > \underline{\pi}$. The agent with the higher output wins the tournament and receives the winner prize W, whereas the agent with the lower output receives the loser prize L. In case that

[5]This parameterized distribution formulation of rank-order tournaments was used recently also by Kräkel and Nieken (2015).

[6]Our assumptions on the effort cost function are slightly different than the assumptions in Gill and Stone (2010), who posit that $c'(0) \geq 0$ and $c'''(e_i) \geq 0$. Furthermore, as becomes apparent below, in our model each contestant's winning probability is linear in efforts, whereas Gill and Stone (2010) allow for a more general form of a contestant's probability to win the tournament. To guarantee a fair comparison of results in the light of these differences, we fully replicate the findings by Gill and Stone (2010) with regard to tournaments with homogeneous contestants.

both agents produce the same output, the winner of the tournament is determined by the flip of a fair coin. Hence, given effort choices e_i and e_j, the probability of agent i receiving the winner prize amounts to

$$P_i(e_i, e_j) \;=\; e_i(1-e_j) + \frac{1}{2}\left[e_i e_j + (1-e_i)(1-e_j)\right] \;=\; \frac{1+e_i-e_j}{2}.$$

Both agents have reference-dependent preferences and are expectation-based loss averse à la Kőszegi and Rabin (2006). Specifically, utility is additively separable across the money dimension and the effort dimension. Furthermore, in each dimension an agent not only experiences standard material utility but also psychological gain-loss utility from comparing the actual consumption outcome to a reference point. This reference point is shaped by the agent's recently held rational expectations; i.e., she compares the material utility of the actual outcome to the material utility of each outcome that she expected to possibly occur, where each such comparison is weighted with the probability that the agent assigned to the respective reference outcome given her recent expectations. Formally, given the agents exert efforts e_i and e_j, respectively, and agent i expected herself to exert effort $\hat{e}_i$, then agent i's expected utility amounts to

$$\begin{aligned} U^i(e_i, \hat{e}_i, e_j) = P_i(e_i, e_j)\left\{u(W) + \eta[1 - P_i(\hat{e}_i, e_j)]\mu(\Delta)\right\} \\ + [1 - P_i(e_i, e_j)]\left\{u(L) - \eta\lambda P_i(\hat{e}_i, e_j)\mu(\Delta)\right\} \\ - c(e_i) + \eta\mu\left(c(\hat{e}_i) - c(e_i)\right). \quad (4.1) \end{aligned}$$

Here, $\eta \geq 0$ is the weight a decision maker attaches to gain-loss utility relative to intrinsic utility and

$$\mu(x) = \begin{cases} x & \text{if } x \geq 0 \\ \lambda x & \text{if } x < 0 \end{cases}$$

is a universal gain-loss function, where $\lambda > 1$ captures loss aversion in the sense that a loss looms larger than an equally sized gain.

As a benchmark, consider the case of loss-neutral contestants for whom $\eta = 0$. Given his opponent's effort choice e_j, according to (4.1) agent i chooses effort e_i to maximize the difference between his expected material utility and his effort cost, $U^i(e_i, \hat{e}_i, e_j) = P_i(e_i, e_j)u(W) + [1 - P_i(e_i, e_j)]u(L) - c(e_i)$, which is strictly concave in e_i. Hence, agent i's best response to agent j exerting effort e_j is characterized by the first-order condition $\partial U^i(e_i, \hat{e}_i, e_j)/\partial e_i = 0$, or, equivalently, $c'(e_i) = \Delta/2$. As this first-order condition does not depend on agent j's effort choice, we conclude the following:

Observation 1. *Suppose $\eta = 0$. The unique Nash equilibrium is symmetric with $(e_A, e_B) = (e^{NE}, e^{NE})$, where e^{NE} satisfies $c'(e^{NE}) = \Delta/2$ and is a strictly dominant strategy for each contestant.*

4.3. Choice Acclimating Personal Equilibrium

In analogy to Gill and Stone (2010), suppose that each agent is called to make her effort choice without having much time to contemplate this decision. The expectation about agent i's own effort choice thus is determined by his actual effort choice, i.e., $\hat{e}_i \equiv e_i$. For a given effort choice e_j of the opponent, expecting to choose and actually choosing effort level e_i yields expected utility

$$U^i(e_i, e_i, e_j) = P(e_i, e_j)[u(W) + \eta(1 - P(e_i, e_j))\Delta] \\ + [1 - P(e_i, e_j)][u(L) - \eta\lambda P(e_i, e_j)\Delta] - c(e_i). \quad (4.2)$$

In the context of individual decision making, Kőszegi and Rabin (2007) define a choice-acclimating personal equilibrium (CPE) as a choice of action that maximizes the decision maker's expected utility given that his expectations correctly reflect the consequences of his action. Just like a Nash equilibrium, an equilibrium of the game under the assumption of choice-acclimating expectations is characterized by mutual best responses, where each agent's best response constitutes a CPE given the opponent's choice of effort.

Definition 4.1. *The effort choices $(\tilde{e}_A, \tilde{e}_B)$ represent a choice-acclimating Nash equilibrium (CPNE) in the rank-order tournament if and only if for all $i \in \{A, B\}$ and $j \in \{A, B\}$ with $j \neq i$,*

$$U^i(\tilde{e}_i, \tilde{e}_i, \tilde{e}_j) \geq U^i(e_i, e_i, \tilde{e}_j) \qquad \forall e_i \in [0, 1]. \quad (4.3)$$

To analyze equilibrium play under choice-acclimating expectations, we begin by establishing some properties of agent i's expected utility function $U^i(e_i, e_i, e_j)$. Differentiation of (4.2) with respect to e_i yields

$$\frac{\partial U^i(e_i, e_i, e_j)}{\partial e_i} = \frac{\Delta}{2}\left[1 + \eta(\lambda - 1)(e_i - e_j)\right] - c'(e_i) \quad (4.4)$$

and

$$\frac{\partial^2 U^i(e_i, e_i, e_j)}{\partial e_i^2} = \frac{\Delta}{2}\eta(\lambda - 1) - c''(e_i). \quad (4.5)$$

Since $c'''(e_i) > 0$, (4.5) implies that $U^i(e_i, e_i, e_j)$ has at most one inflection point. Furthermore, as $\lim_{e_i \to 1} c'(e_i) = \infty$ implies $\lim_{e_i \to 1} c''(e_i) = \infty$, from (4.4) and (4.5) it follows that $U^i(e_i, e_i, e_j)$ is strictly decreasing and concave as e_i becomes sufficiently close to unity. In consequence, $U^i(e_i, e_i, e_j)$ is either concave for all $e_i \in [0, 1]$ or, if an inflection point exists, convex (concave) for values of e_i below (above) the inflection point.

These observations have direct implications regarding the effort levels that are candidates for being agent i's best response to effort level e_j of his opponent. First, maximum effort $e_i = 1$ never is a best response. Second, $U^i(e_i, e_i, e_j)$ has at most one local maximum in the open interval $(0, 1)$. We denote this local maximizer (if it exists) by $e_i^*(e_j)$, which solves the first-order condition $\partial U^i(e_i^*(e_j), e_i^*(e_j), e_j)/\partial e_i = 0$ or, equivalently,

$$c'(e_i^*(e_j)) = \frac{\Delta}{2}\left[1 + \eta(\lambda - 1)(e_i^*(e_j) - e_j)\right]. \tag{4.6}$$

As $e_i^*(e_j)$ is a local maximizer, it is readily established that $de_i^*(e_j)/de_j < 0$ and $d^2 e_i^*(e_j)/de_j^2 < 0$. Thus, whenever agent i's best response to agent j's effort is a strictly positive amount of effort, a slight increase in agent j's effort decreases agent i's best response—i.e., effort choices are strategic substitutes. Third, and finally, the boundary effort choice $e_i = 0$ always represents a candidate for a local extreme point of $U^i(e_i, e_i, e_j)$. Hence, we have two candidates for agent i's best response to effort level e_j: the interior local maximizer $e_i = e_i^*(e_j)$ or minimum effort $e_i = 0$. With

$$\frac{\partial\left[U^i(0, 0, e_j) - U^i(e_i^*(e_j), e_i^*(e_j), e_j)\right]}{\partial e_j} = \frac{\Delta}{2}\eta(\lambda - 1)e_i^*(e_j) > 0, \tag{4.7}$$

exerting zero effort becomes more attractive relative to choosing the interior local maximizer $e_i^*(e_j)$ as e_j increases. In consequence, we should not be surprised to see minimum effort being agent i's best response in particular for high effort levels of his opponent. For low levels of the opponent's effort, in contrast, agent i's best response always involves the interior effort level $e_i^*(e_j)$ because $\partial U^i(0, 0, e_j)/\partial e_i > 0$.

Lemma 1. *Either agent i's best response function under choice-acclimating expectations is given by $BR_i(e_j) = e_i^*(e_j)$ for all $e_j \in [0, 1]$, or there exists $\bar{e} \in (0, 1]$ such that agent i's best response is given by*

$$BR_i(e_j) = \begin{cases} e_i^*(e_j) & \text{if } e_j \leq \bar{e} \\ 0 & \text{if } e_j \geq \bar{e} \end{cases},$$

where $e_i^*(e_j) \in (0, 1)$, $\frac{de_i^*(e_j)}{de_j} < 0$ and $\frac{d^2 e_i^*(e_j)}{de_j{}^2} < 0$.

While we defer the details of the proof of Lemma 1 to the appendix, the qualitative features of the best-response function under choice-acclimating beliefs can be summarized as follows. First, for $\eta(\lambda - 1) < 1$, a player attaches greater weight to expected material utility than to the expected net loss.[7] In this case, exerting no effort at all can never be optimal, such that agent i's best response to e_j is always given by the interior local maximizer $e_i^*(e_j)$ (see the left panel of Figure 4.1). With psychological concerns being of little importance, slightly increasing effort above zero improves expected utility due to its positive impact on expected material utility. Second, if $\eta(\lambda - 1) \geq 1$, psychological concerns are more important than material concerns.

[7] Herweg, Müller, and Weinschenk (2010) refer to the case where $\eta(\lambda - 1) < 1$ as "no dominance of gain-loss utility".

In this case, total resignation may actually be a best response to high effort levels of the opponent. By exerting minimum effort a contestant moderates his expectation about the tournament outcome, which yields a lower expected net loss and, as an added benefit, minimizes effort cost. More specifically, if effort costs are strongly convex—i.e., $\frac{2}{\Delta}c''(0) > \eta(\lambda - 1)$—expected utility is strictly concave for all e_j and the best response of agent i decreases continuously to zero (see the middle panel of Figure 4.1). If, on the other hand, loss aversion is rather strong—i.e., $\eta(\lambda - 1) \geq \frac{2}{\Delta}c''(0)$—agent i's best response may display a discontinuity (see the right panel of Figure 4.1). Intuitively, with strong loss aversion, the concern to avoid losses is very prevalent and each agent's main objective is to reduce the uncertainty regarding the tournament outcome. With regard to only psychological gain-loss utility, if the opponent exerts little effort, a contestant reduces the likelihood to experience a loss by exerting maximum effort, which yields an almost certain win of the tournament. However, as the opponent's effort—and thus, ceteris paribus, his winning probability—increases, there is a threshold for the opponent's effort above which uncertainty about the tournament outcome is not minimized by agent i exerting maximum effort but by exerting minimum effort. This discontinuity in the level of effort that minimizes uncertainty in the tournament translates into a downward discontinuity in agent i's best response function.

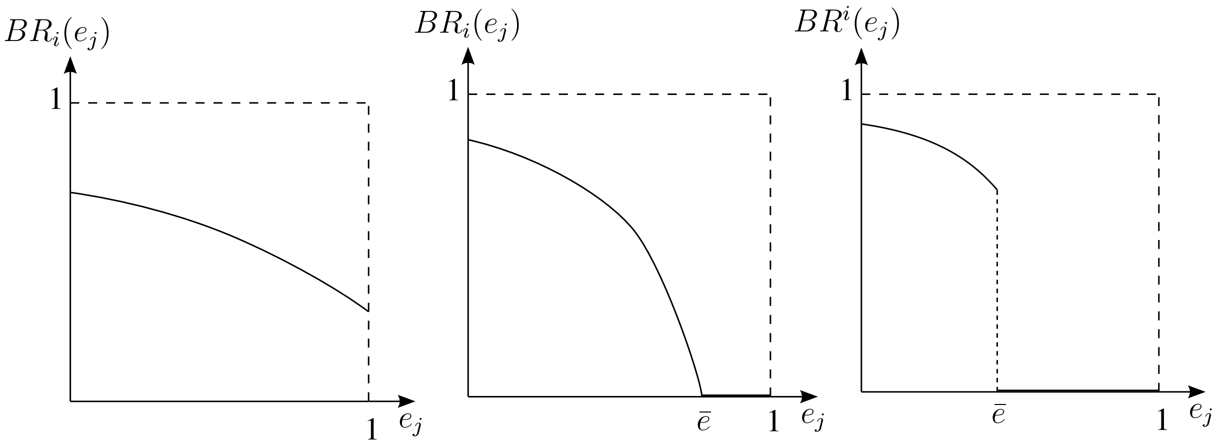

Figure 4.1.: Player i's best response function for different degrees of loss aversion.

Based on these insights, we can now re-establish the result of Gill and Stone (2010) with regard to asymptotic stability and nonexistence of symmetric equilibria in our parameterized distribution formulation of a rank-order tournament.

Proposition 4.1. *Any symmetric CPNE must be identical to the unique Nash equilibrium* (e^{NE}, e^{NE}).

(i) For $\eta(\lambda - 1) \in \left[\frac{c''[c'^{-1}(\frac{\Delta}{2})]}{\Delta}, \frac{2c''[c'^{-1}(\frac{\Delta}{2})]}{\Delta}\right]$*, such a symmetric CPNE will be asymptotically unstable.*

(ii) For $\eta(\lambda - 1) > \frac{2c''[c'^{-1}(\frac{\Delta}{2})]}{\Delta}$*, such a symmetric CPNE cannot exist.*

With agent i's best response being decreasing in agent j's effort level, there can exist at most one symmetric CPNE. Note that psychological concerns are effectively absent from marginal utility in Equation (4.4) if both agents exert identical effort, such that in a symmetric CPNE each agent's best response balances marginal effort cost with marginal expected money income. In consequence, given a symmetric CPNE exists, it necessarily coincides with the unique Nash equilibrium (e^{NE}, e^{NE}). For the symmetric CPNE to be asymptotically stable, the best response function of player i has to be sufficiently flat at the Nash equilibrium, which is the case if and only if $\eta(\lambda - 1) < \frac{c''[c'^{-1}(\frac{\Delta}{2})]}{\Delta}$. For larger values of $\eta(\lambda - 1)$, the best response function of an agent is sufficiently steep such that the symmetric CPNE becomes asymptotically unstable. Finally, the symmetric equilibrium ceases to exist if and only if agent i's best-response function discontinuously drops to zero before crossing the 45°-line. This definitely is the case if $\eta(\lambda - 1) > \frac{2c''[c'^{-1}(\frac{\Delta}{2})]}{\Delta}$, which yields that playing $e_i = e^{NE}$ no longer constitutes a local interior maximizer but a local minimizer of $U^i(e_i, e_i, e^{NE})$.

Whenever a symmetric CPNE does not exist, there exist two asymmetric CPNEs in which exactly one agent exerts strictly positive effort whereas the other agent resigns and does not exert any effort.

Proposition 4.2. *For $\eta(\lambda - 1)$ sufficiently large, two asymmetric CPNEs characterized by $(e_A^{CPNE}, e_B^{CPNE}) = (e^*(0), 0)$ and $(e_A^{CPNE}, e_B^{CPNE}) = (0, e^*(0))$ with $e^*(0) > 0$ exist. These asymmetric CPNEs are asymptotically stable.*

As noted before, Propositions 4.1 and 4.2 replicate the results provided by Gill and Stone (2010) in our parameterized distribution formulation of the tournament environment. Furthermore, the slightly more tractable approach allows us to complement these important findings. In particular, we can specify the exact degree of loss aversion that is necessary and sufficient for the symmetric equilibrium to cease to exist.

Proposition 4.3. *There exists $\overline{\chi} \in \left(1, \frac{2c''(c'^{-1}(\frac{\Delta}{2}))}{\Delta}\right)$ such that the symmetric CPNE exists if and only if $\eta(\lambda - 1) \leq \overline{\chi}$.*

Proposition 4.3 demonstrates that a symmetric equilibrium exists for moderate degrees of loss aversion.[8] In fact, the contestants' concern for psychological gain-loss utility must outweigh their concern for consumption utility for the symmetric equilibrium not to exist. Such a strong degree of loss aversion, however, would also imply violations of stochastic dominance (Kőszegi and Rabin, 2007). Note that, referring to Proposition 4.1 (i), the symmetric CPNE may already be asymptotically unstable as it ceases to exist. We thus conclude that a focus

[8] As shown by Gill and Stone (2010), this picture changes if contestants are heterogeneous with respect to their ability, in which case all CPNEs are asymmetric.

on symmetric equilibria in rank-order tournaments may be justifiable if players have choice-acclimating expectations and and minimal efforts come without costs.[9]

4.4. Personal Equilibrium

In rank-order tournaments, it is sometimes conceivable to think of agents having time to ponder their behavior in the tournament before it actually takes place. For example, teams or single players may often enter the tournament with a given "game plan", political candidates may elaborate on their campaign months before it takes place, and workers may try to elicit bits and pieces of their peers' strategies for the upcoming promotion tournament to adapt their own strategy accordingly. In this case agents enter the tournament with fixed expectations about their behavior. The assumption of fixed expectations is particularly compelling for later stages in a multi-stage tournament. Sure enough contestants form expectations about their chances to win the tournament and the resulting necessary efforts latest when they enter the first stage. As a direct consequence, they begin every following stage with fixed expectations.

Formally, for the case of fixed expectations, we assume that agent i makes her actual effort choice e_i for expectations $\hat{e}_i$ regarding her own behavior. To guarantee internal consistency of expectations and actual behavior, we apply the concept of personal equilibrium (PE) as defined in Kőszegi and Rabin (2006); i.e., we require that a person can reasonably expect a particular course of action only if she is willing to follow it through given her expectations. The following definition extends this idea of internal consistency to the outcome of tournament play.

Definition 4.2. *The effort choices $(\tilde{e}_A, \tilde{e}_B)$ represent a personal Nash equilibrium (PNE) in the rank-order tournament if and only if for all $i \in \{A, B\}$ and $j \in \{A, B\}$ with $j \neq i$,*

$$U^i(\tilde{e}_i, \tilde{e}_i, \tilde{e}_j) \geq U^i(e_i, \tilde{e}_i, \tilde{e}_j) \qquad \forall e_i \in [0, 1]. \tag{4.8}$$

Essentially, in a PNE each agent's effort choice constitutes a PE given her opponent's effort choice. In order to identify the set of PNEs in the rank-order tournament, we begin by analyzing the set of PEs for agent i for a given effort e_j of her opponent. To this end, note that a necessary condition for effort level $\tilde{e}_i$ to be a PE is that neither a marginal upward deviation nor a marginal downward deviation is strictly profitable for agent i. Formally,

$$\left.\frac{\partial U^i(e_i, \tilde{e}_i, e_j)}{\partial e_i}\right|_{e_i \searrow \tilde{e}_i} = \frac{\Delta}{2}\left\{1 + \eta\left[1 + (\lambda - 1)P(\tilde{e}_i, e_j)\right]\right\} - (1 + \eta\lambda)c'(\tilde{e}_i) \leq 0 \tag{4.9}$$

[9] There exist alternative specifications to the cost of effort function such that the symmetric CPNE ceases to exist for a critical value of $\eta(\lambda - 1) \leq 1$. As suggested by the linear example in Gill and Stone (2010), this holds true if exerting a minimum amount of effort is sufficiently costly and $c''(\cdot)$ is small.

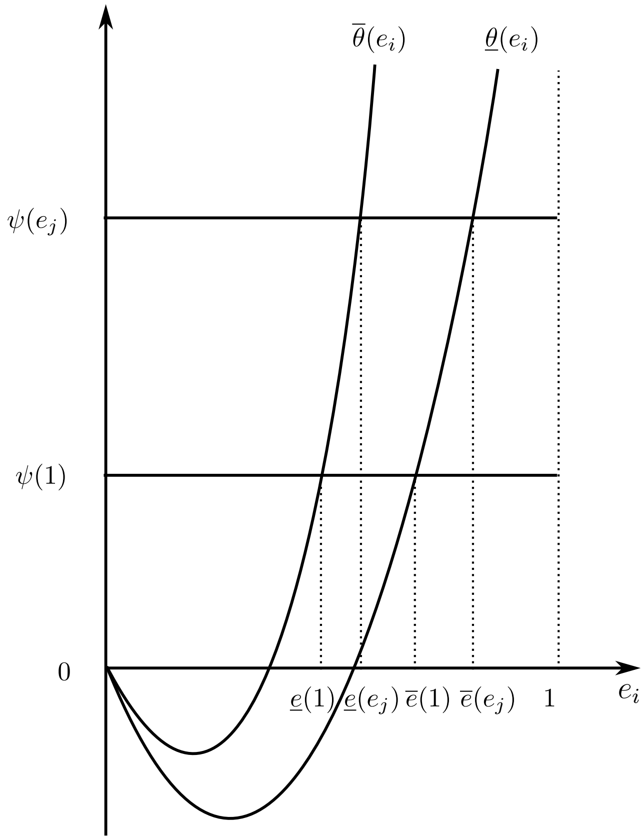

Figure 4.2.: Construction of the set $\Theta_i^{PE}(e_j)$.

and

$$\left.\frac{\partial U^i(e_i, \tilde{e}_i, e_j)}{\partial e_i}\right|_{e_i \nearrow \tilde{e}_i} = \frac{\Delta}{2}\left\{1 + \eta\left[1 + (\lambda - 1)P(\tilde{e}_i, e_j)\right]\right\} - (1 + \eta)c'(\tilde{e}_i) \geq 0, \tag{4.10}$$

have to hold simultaneously. Here, $\frac{\partial^2 U^i(e_i,\tilde{e}_i,e_j)}{\partial e_i^2} = -(1 + \eta)c''(e_i) < 0$ for all $e_i < \tilde{e}_i$ and $\frac{\partial^2 U^i(e_i,\tilde{e}_i,e_j)}{\partial e_i^2} = -(1 + \eta\lambda)c''(e_i) < 0$ for all $e_i > \tilde{e}_i$. Thus, given (4.9) and (4.10) are satisfied, the expected utility of player i is strictly increasing in e_i for $e_i < \tilde{e}_i$ and strictly decreasing for $e_i > \tilde{e}_i$. Hence, (4.9) and (4.10) together constitute not only a necessary but also a sufficient condition for $\tilde{e}_i$ to be a PE. For agent i, we denote the resulting set of PEs for a given effort choice e_j of her opponents by

$$\Theta_i^{PE}(e_j) = \left\{\tilde{e}_i \in [0, 1] \mid \text{(4.9) and (4.10) are satisfied}\right\}. \tag{4.11}$$

In order to characterize this set, define the functions

$$\underline{\theta}(\tilde{e}_i) \equiv 2c'(\tilde{e}_i)(1 + \eta) - \Delta\frac{\eta(\lambda - 1)\tilde{e}_i}{2}, \tag{4.12}$$

$$\bar{\theta}(\tilde{e}_i) \equiv 2c'(\tilde{e}_i)(1 + \eta\lambda) - \Delta\frac{\eta(\lambda - 1)\tilde{e}_i}{2}, \tag{4.13}$$

and

$$\psi(e_j) \equiv \Delta\left[1+\eta+\frac{\eta(\lambda-1)}{2}(1-e_j)\right], \tag{4.14}$$

which are illustrated in Figure 4.2. Condition (4.9) can be rewritten as $\bar{\theta}(\tilde{e}_i) \geq \psi(e_j)$. Note that $\psi(e_j)$ is strictly decreasing and that $\psi(1) > 0$ such that $\psi(e_j)$ is strictly positive for all $e_j \in [0,1]$. Next, consider the function $\bar{\theta}(\tilde{e}_i)$. By our assumptions on the effort cost function, we have $\bar{\theta}(0) = 0$ and $\lim_{e_i \to 1} \bar{\theta}(e_i) = \infty$. Hence, the intermediate value theorem guarantees that there exists $\underline{e}_i(e_j) \in (0,1)$ such that $\bar{\theta}(\underline{e}_i(e_j)) = \psi(e_j)$. Due to the strict convexity of $\bar{\theta}(e_i)$, $\underline{e}_i(e_j)$ is uniquely determined and effort levels below $\underline{e}_i(e_j)$ do not constitute a PE for agent i given her opponent exerts effort e_j. Similarly, condition (4.10) can be rewritten as $\underline{\theta}(\tilde{e}_i) \leq \psi(e_j)$. By analogous reasoning, we can establish the existence of $\bar{e}_i(e_j) \in (0,1)$ such that $\underline{\theta}(\bar{e}_i(e_j)) = \psi(e_j)$ and effort levels above $\bar{e}_i(e_j)$ do not constitute a PE for agent i, either. Finally, since $\bar{\theta}(e_i) < \underline{\theta}(e_i)$ for all $e_i \in (0,1]$, we have $\underline{e}_i(e_j) < \bar{e}_i(e_j)$. This allows us to establish the following observation.

Lemma 2. *Given $e_j \in [0,1]$, $\Theta_i^{PE}(e_j) = [\underline{e}_i(e_j), \bar{e}_i(e_j)] \subseteq (0,1)$. Furthermore, $\underline{e}_i(e_j)$ and $\bar{e}_i(e_j)$ are continuous, strictly decreasing, and strictly concave.*

According to Lemma 2, agent i can credibly expect only to exert a moderate effort level herself. For any fixed expectation, increasing the effort beyond this expectation involves a tradeoff for the agent. On the one hand, an increase in effort improves her chances to win the tournament and to experience a gain and at the same time it reduces the probability to obtain the loser prize and to experience a loss. On the other hand, the corresponding increase in effort implies higher effort costs and leads to a certain loss in the effort-cost dimension. Due to this tradeoff, expecting to exert a fairly low effort level, $e_i < \underline{e}_i(e_j)$ is not a credible plan for agent i. In this case, the convexity of the effort cost function implies that the latter drawback is rather small and more than outweighed by the former benefit, such that a deviation to a higher effort level is profitable. Likewise, expecting to exert a fairly high effort level, $e_i > \bar{e}_i(e_j)$, neither is a credible plan for agent i. In this case, the benefit of reducing effort costs by decreasing effort beyond this expectation more than outweighs the drawbacks associated with the decrease in agent i's winning probability. Agent A's set of personal equilibria in dependence of agent B's effort choice is depicted in the left panel of Figure 4.3.

Finally, we use the preceding characterization of an agent's set of PEs to derive equilibrium behavior. According to Definition 4.2, the set of PNEs is given by

$$\Theta^{PNE} = \{(e_A, e_B) \in [0,1]^2 \mid e_A \in \Theta_A^{PE}(e_B) \text{ and } e_B \in \Theta_B^{PE}(e_A)\}. \tag{4.15}$$

By the properties of the agents' PE correspondences listed in Lemma 2, it follows immediately that there always exists a PNE. Furthermore, as becomes apparent from the right panel of Figure

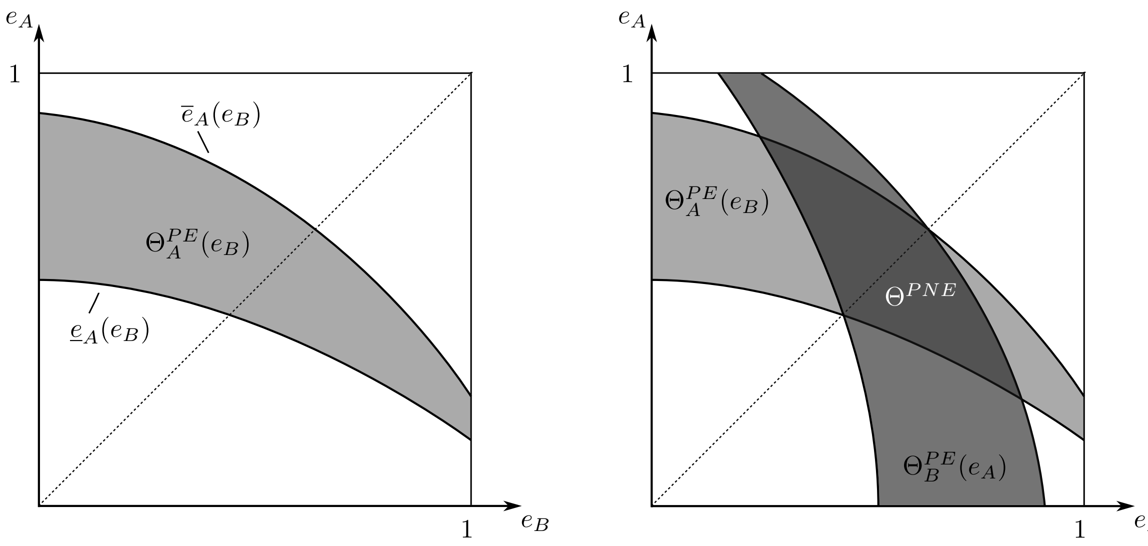

Figure 4.3.: The left panel depicts the correspondence of PEs for Agent A. The right panel the resulting set of PNEs for the tournament.

4.3, which depicts the set of PNEs, next to asymmetric PNEs there always—i.e., for any degree of loss aversion—exist symmetric PNEs in which both agents exert the same level of effort.

Proposition 4.4. *There exist symmetric PNEs in which the agents exert a moderate level of effort; i.e., there exist $\underline{e}$ and $\bar{e}$ with $0 < \underline{e} < \bar{e} < 1$, such that $(e, e) \in \Theta^{PNE}$ for all $e \in [\underline{e}, \bar{e}]$.*

The existence of symmetric PNEs for all degrees of loss aversion distinguishes the case of fixed expectations from the case of choice-acclimating expectations. As stated in Kőszegi and Rabin (2007), choice-acclimating expectations result in stronger risk aversion than fixed expectations. For the case of a rank-order tournament, this induces contestants with choice-acclimating expectations to dislike the uncertainty in a symmetric equilibrium so intensely that they choose rather different effort levels, with one agent completely resigning and exerting no effort at all. By resigning this agent reduces his chances to win the tournament but at the same time he is able to moderate her expectations and thus to dampen the pain of a potential loss. With fixed expectations, in contrast, exerting an identical, moderate level of effort is a credible plan for both agents—irrespective of their degree of loss aversion. Here, when expecting to exert moderate effort, resignation would reduce an agent's chances to win the tournament without moderating his expectations, such that resignation would badly disappoint the agent's hopes of winning the tournament. This inevitable increase in the likelihood to experience a loss makes the potential deviation unattractive and a symmetric equilibrium always exists if expectations are fixed.

4.5. Preferred Personal Nash Equilibrium

In the previous sections, we showed that expectation-based loss aversion per se does not necessarily lead to asymmetric equilibria. In particular, for the case of fixed expectations there always exist both symmetric and asymmetric equilibria. The multiplicity of equilibria raises the question which of the prevalent equilibria is most suitable to describe the contestants' behavior. To answer this question in the context of individual decision making, Kőszegi and Rabin (2007) propose the notion of preferred personal equilibrium (PPE) as an equilibrium refinement. The PPE is the PE that promises the highest expected utility among all PEs. If players are able to select their most preferred personal equilibrium given any strategy of the opponent, this concept can also be adopted for the context of strategic interaction. We define a preferred personal Nash equilibrium (PPNE) such that every player plays a PPE given his opponent's strategy.

Definition 4.3. *The effort choices* $(\tilde{e}_A, \tilde{e}_B)$ *represent a preferred personal Nash equilibrium (PPNE) in the rank-order tournament if and only if for all* $i \in \{A, B\}$ *and* $j \in \{A, B\}$ *with* $j \neq i$, $\tilde{e}_i \in \Theta_i^{PE}(\tilde{e}_j)$ *and*

$$U^i(\tilde{e}_i, \tilde{e}_i, \tilde{e}_j) \geq U^i(e_i, e_i, \tilde{e}_j) \qquad \forall e_i \in \Theta_i^{PE}(\tilde{e}_j). \tag{4.16}$$

Recall that the CPE is an agent's most profitable action among all his actions provided that his expectations are consistent with consequences of the action he actually takes. Hence, in the context of individual decision making, if a CPE constitutes a PE for a player it also is a PPE. By the same reasoning, a CPNE that constitutes a PNE is also a PPNE. In Section 4.4, we showed that any symmetric CPNE indeed constitutes a PNE and thus is a PPNE. Proposition 4.3 then allows us to conclude that there always exists a symmetric PPNE as long as gain loss utility does not dominate material utility. The persistence of the symmetric PPNE, however, is even stronger than that of the symmetric CPNE: even if loss aversion dominates material utility so intensely that the symmetric CPNE ceases to exist, there may still exist a symmetric PPNE.

Proposition 4.5. *There exists* $\tilde{\chi} > \overline{\chi}$ *such that for all* $\eta(\lambda - 1) < \tilde{\chi}$ *the symmetric Nash equilibrium is a PPNE.*

The stronger persistence of the symmetric PPNE arises because a contestant with fixed expectations is more limited in his choice of effort—namely to those effort levels that constitute a PE—than a contestant with choice-acclimating expectations. As explained in Section 4.3, the symmetric CPNE ceases to exist for high degrees of loss aversion, because one player ultimately chooses to resign for the purpose of reducing uncertainty in the tournament. As was established in Section 4.4, however, for fixed expectations not exerting any effort is never a credible plan. If a player expected not to exert any effort at all, he would always be better off by surprising himself and exerting slightly positive effort, which comes without cost (by $c'(0) = 0$) but strictly increases his chances of winning. Thus, a contestant who is restricted to choose an effort level that constitutes a PE cannot reduce uncertainty to the same extent as a contestant who is not restricted in this regard. In consequence, possible deviations from the symmetric equilibrium are less attractive and the symmetric Nash equilibrium is a PPNE for even stronger degrees of loss aversion than for which it is a CPNE.

We conclude that, if players enter the tournament with a game plan, symmetric equilibria are quite persistent. In particular, they always exist if players are not able to select among their PEs. Moreover, even if players have this ability, there always exists a symmetric PPNE as long as players' concerns for gain loss utility does not clearly outweigh those for material utility. In this case, under PNE and PPNE the behavior of expectation-based loss-averse players resembles the behavior of players with standard utility or players with exogenously given reference points (cf. Gill and Stone (2010) and Gill and Prowse (2012)).

4.6. Conclusion

Many of our insights about rank-order tournaments build upon the premise that symmetric equilibria exist. As shown by Gill and Stone (2010), the existence of symmetric equilibria may fail if contestants are expectation-based loss averse and have choice-acclimating expectations. In this note, we complement their work by delineating the circumstances under which a focus on symmetric equilibria is nevertheless justifiable even if players are expectation-based loss averse. First, if contestants' concerns for psychological gain-loss utility do not outweigh material consumption utility and a minimal effort comes without costs, the existence of a symmetric equilibrium is guaranteed. Second, symmetric equilibria also exist for all degrees of loss aversion if players enter the tournament with fixed expectations. Third, while for fixed expectations also asymmetric equilibria exist, the symmetric Nash equilibrium always prevails if each player follows his preferred credible game plan and concerns for psychological gain-loss utility do not outweigh material consumption utility.

This note also adds to the emerging literature that analyzes strategic interaction of expectation-based loss-averse agents by investigating how the equilibrium concepts of Nash equilibrium, personal Nash equilibrium, and choice-acclimating Nash equilibrium relate to each other.[10] Regarding rank-order tournaments, a desirable next step would be to explore the implications of expectation-based loss aversion in dynamic tournaments à la Rosen (1986), where choice-acclimating expectations and lagged fixed expectations do not necessarily represent alternative

[10]Strategic interaction of expectation-based loss-averse agents has been primarily analyzed in rather specific environments like tournaments (Gill and Stone, 2010; Bergerhoff and Vosen, 2014), team production (Gill and Stone, 2015), team compensation (Daido and Murooka, 2014), or auctions (Lange and Ratan, 2010). A more general approach is presented in Dato, Grunewald, and Müller (2015b).

modeling choices: as tournament play evolves, choice-acclimating expectations might apply in the very first round, whereas decisions in later rounds are taken with a fixed set of expectations.

A. Appendices

A.1. Appendix to Chapter I

Proof of Lemma 1.1: Following the derivations in the main text, the aggregate profit of the colluding firms is equal to:

$$\pi^M(p,\hat{p}) = [\alpha\,(1-F(p)) + (1-\alpha)(1-F(p+\hat{p}))](p+\hat{p}).$$

To derive the desired result we need to make two case distinctions:

- Suppose that $p+\hat{p} \leq \overline{v}$.
  Then it holds that some sophisticates still participate in the market, i.e. $1-F(p+\hat{p}) > 0$. Now we want to derive the optimal base good price:

  $$\frac{\partial \pi^M}{\partial p} = [-\alpha f(p) - (1-\alpha) f(p+\hat{p})](p+\hat{p}) + \alpha(1-F(p)) + (1-\alpha)(1-F(p+\hat{p})) \overset{!}{=} 0$$
  $$\Leftrightarrow\; p = \frac{1}{2}(\overline{v} - (2-\alpha)\hat{p}). \tag{A.1}$$

  Recall that we have assumed that some sophisticates participate in the market. Then, if the add-on price is below its maximum value, the firms can increase their profit by lowering the base good price marginally while increasing the add-on price by the same amount. This leaves the total price of the product bundle and the demand from sophisticates unchanged but increases the demand from myopic consumers. Hence, it is profitable to increase the add-on price $\hat{p}$ until it reaches its upper bound $\bar{p}$ or until the base good price reaches its lower bound. If the add-on price reaches its upper bound, the value of the base good price (A.1) is:

  $$p = \frac{1}{2}(\overline{v} - (2-\alpha)\bar{p}).$$

  Since we have assumed that $\overline{v} \leq \bar{p}$, the above expression would yield a negative price of the base good, i.e. $p < 0$. This is not allowed in terms of the model due to the lower bound for the base good. Hence, the optimal base good price is $p^M = 0$.

- Suppose that $p+\hat{p} > \overline{v}$.
  This implies that no sophisticated consumer participates in the market, i.e. $1-$

$F(p+\hat{p}) = 0$. Note here that in this case the profit is strictly increasing in the add-on price $\hat{p}$. Hence, firms will set the add-on price equal to the maximum add-on price, i.e. $\hat{p} = \bar{p}$. Maximizing the profit of firms over the base good price then yields:

$$\begin{aligned} \frac{\partial \pi^M}{\partial p} &= -\alpha f(p)(p+\bar{p}) + \alpha(1-F(p)) \stackrel{!}{=} 0 \\ \Leftrightarrow \quad p &= \frac{\overline{v}}{2}\left(1 - \frac{\bar{p}}{\overline{v}}\right). \end{aligned}$$

This base good price is negative if $\overline{v} < \bar{p}$, as was assumed in Assumption 1. Hence, due to the lower bound for the base good price, it is again optimal to set $p^M = 0$.

□

Proof of Proposition 1.2: To prove that an optimal deviation from collusion includes an unshrouding of the add-on, we have to compare the different possibilities how firms can collude and how a firm could deviate from the collusion.

- First consider the case in which the deviating firm decides to unshroud the add-on. As we have already shown in the main text, the deviating firm will set an add-on price of $\hat{p}_i^{dev} = \frac{\overline{v}}{2}$ and earn a deviation profit of $\pi_i^{dev} = \frac{\overline{v}}{4}$. The add-on price $\hat{p}_i^{dev} = \frac{\overline{v}}{2}$ is feasible since $\overline{v} \leq \bar{p}$, which ensures that $\hat{p}_i^{dev} \leq \bar{p}$.
- Now suppose that the deviating firm does not unshroud the add-on. In this case, we have to distinguish whether the deviating firm charges an add-on price below $\overline{v}$ (inner solution) or an add-on price above $\overline{v}$ (corner solution).
  - If the deviating firm charges an add-on price below $\overline{v}$, the optimal add-on price was $\pi_i^{dev} = \left[\frac{\alpha}{n} + (1-\alpha)(1-\frac{\hat{p}}{\overline{v}})\right]\hat{p}$, yielding a deviation profit of $\pi_i^{dev} = \frac{\overline{v}}{4(1-\alpha)}\left(1-\alpha\frac{n-1}{n}\right)^2$. Note that, for some parameter constellations, it might be the case that $\hat{p}_i^{dev} = \frac{\overline{v}}{2(1-\alpha)}\left(1-\alpha\frac{n-1}{n}\right) > \overline{v}$. In this case, the derived add-on price that a deviating firm sets is not an inner solution since no sophisticated consumer will buy the product bundle. It holds that the add-on price of the optimal deviation is feasible if:

    $$\begin{aligned} \hat{p}_i^{dev} = &\frac{\overline{v}}{2(1-\alpha)}\left(1-\alpha\frac{n-1}{n}\right) \leq \overline{v} \\ \Leftrightarrow \alpha \quad &\leq \frac{n}{n+1}. \end{aligned}$$

    Hence, the deviation add-on price of the inner solution $\pi_i^{dev} = \left[\frac{\alpha}{n} + (1-\alpha)(1-\frac{\hat{p}}{\overline{v}})\right]\hat{p}$ is only feasible if $\alpha \leq \frac{n}{n+1}$.
  - Now suppose that the deviating firm charges an add-on price above $\overline{v}$. In this case, the profit function of the deviating firm is:

    $$\pi_i^{dev} = \left[\frac{\alpha}{n}\right]\hat{p}.$$

It is easy to verify that the add-on price that maximizes the above profit function is $\hat{p} = \bar{p}$, which would yield a profit of $\pi_i^{dev} = \frac{1}{n}\alpha\bar{p}$. As already argued in the main text, this cannot be a profitable deviation since this deviation profit is lower than the profit under collusion. Hence, it will never be a profitable deviation to charge an add-on price above $\overline{v}$.

In summary, we can conclude that a profitably deviating firm that decides not to unshroud the add-on cannot do better than obtaining a profit of $\pi_i^{dev} = \frac{\overline{v}}{4(1-\alpha)}\left(1-\alpha\frac{n-1}{n}\right)^2$. However, recall that the corresponding add-on price is only feasible if $\alpha \leq \frac{n}{n+1}$.

Now we can check what the optimal deviation strategy looks like. Comparing the deviation profit that a firm can obtain by unshrouding to the one without unshrouding yields:

$$\begin{aligned}
&\pi_{i,shrouding}^{dev} &&\leq \pi_{i,unshrouding}^{dev} \\
\Leftrightarrow\ &\frac{\overline{v}}{4(1-\alpha)}\left(1-\alpha\frac{n-1}{n}\right)^2 &&\leq \frac{\overline{v}}{4} \\
\Leftrightarrow\ &\left(1-\alpha\frac{n-1}{n}\right)^2 &&\leq 1-\alpha \\
\Leftrightarrow\ &\frac{2-n}{n}+\alpha\left(\frac{n-1}{n}\right)^2 &&\leq 0 \\
\Leftrightarrow\ &\alpha &&\leq \frac{n(n-2)}{(n-1)^2}.
\end{aligned} \tag{A.2}$$

Recall that a deviation without unshrouding the add-on can only be profitable if $\alpha \leq \frac{n}{n+1}$. Hence, when we want to check whether a deviation without unshrouding can be optimal, we can focus on cases with $\alpha \leq \frac{n}{n+1}$. Then the inequality (A.2) holds if the following relation is fulfilled:

$$\begin{aligned}
\frac{n(n-2)}{(n-1)^2} &\geq \frac{n}{n+1} \\
\Leftrightarrow n &\geq 3,
\end{aligned}$$

which is fulfilled by assumption. Therefore, we can conclude that if a profitable deviation exists, unshrouding will be part of the optimal deviation strategy. □

Proof of Proposition 1.4: Suppose firms coordinate on prices $\hat{p}^{coll}$, yielding some collusive profit π^{coll}: In analogy to the above analysis, the critical discount factor is then given by

$$\delta^* = \frac{n-\frac{\pi^{coll}}{\pi^{dev}}}{n},$$

which is decreasing in the collusion-to-deviation profit ratio $\frac{\pi^{coll}}{\pi^{dev}}$. Hence, the less attractive a deviation is relative to the collusive play, the more stable collusion is.

Recall that if a profitable deviation exists, it is optimal to set $\hat{p}^{dev} = \frac{\overline{v}}{2}$ and unshroud the add-on, thereby earning $\pi_i^{dev} = \frac{\overline{v}}{4}$. This deviation profit is feasible only if firms coordinated on prices $\hat{p}^{coll} \geq \frac{\overline{v}}{2}$ in the collusive play. Then the critical discount factor is minimized by coordinating on maximum profits, which are $\pi^M = \frac{\overline{v}}{4(1-\alpha)}$ or $\pi^M = \alpha\bar{p}$. Hence, coordinating on lower profits cannot stabilize collusion.

Now suppose firms coordinate on prices $\hat{p}^{coll} < \frac{\overline{v}}{2}$, which yields an aggregate profit of $\pi(\hat{p}) = [\alpha + (1-\alpha)(1 - \frac{\hat{p}}{\overline{v}})]\hat{p}$. Playing $\hat{p}^{dev} = \frac{\overline{v}}{2}$ does not correspond to an undercutting anymore and a firm would then optimally deviate by undercutting the collusive price marginally, irrespective of the (un)shrouding decision. This actually follows from the fact that deviation profits are increasing in $\hat{p}^{dev}$ for all prices $\hat{p}^{dev} < \frac{\overline{v}}{2}$.

We will now show that a deviating firm still optimally decides to unshroud the add-on, thereby making use of the fact that the optimal deviation price does not depend on the (un)shrouding decision:

$$\begin{aligned}
\pi_{i,unshrouding}^{dev} &> \pi_{i,shrouding}^{dev} \\
\Leftrightarrow \left[1 - \frac{\hat{p}_i^{dev}}{\overline{v}}\right]\hat{p}_i^{dev} &> \left[\frac{\alpha}{n} + (1-\alpha)(1 - \frac{\hat{p}_i^{dev}}{\overline{v}})\right]\hat{p}_i^{dev} \\
\Leftrightarrow 1 - \frac{\hat{p}^{dev}}{\overline{v}} &> \frac{1}{n}
\end{aligned}$$

This holds since $n \geq 3$ and $\hat{p}^{dev} \leq \frac{\overline{v}}{2}$. Now we that know that a deviating firm will optimally unshroud the add-on and undercut the collusive price marginally, it remains to check whether the critical discount factor can be lowered by coordinating on prices $p^{coll} < \frac{\overline{v}}{2}$. Since colluding with the monopoly price leads to a collusion-to-deviation profit ratio of at least $1/(1-\alpha)$, it must hold that

$$\begin{aligned}
\frac{\pi^{coll}}{\pi^{dev}} = \frac{[\alpha + (1-\alpha)(1 - \frac{\hat{p}}{\overline{v}})]\hat{p}}{\left[1 - \frac{\hat{p}}{\overline{v}}\right]\hat{p}} &> \frac{1}{1-\alpha} \\
\Leftrightarrow (1-\alpha)\left[1 - (1-\alpha)\frac{\hat{p}}{\overline{v}}\right] &> 1 - \frac{\hat{p}}{\overline{v}} \\
\Leftrightarrow \left[1 - (1-\alpha)^2\right]\frac{\hat{p}}{\overline{v}} &> \alpha \\
\Leftrightarrow \alpha\,(2-\alpha)\frac{\hat{p}}{\overline{v}} &> \alpha \\
\Leftrightarrow \hat{p} &> \frac{\overline{v}}{2-\alpha},
\end{aligned}$$

contradicting the assumption that firms colluded with prices $\hat{p}^{coll} < \frac{\overline{v}}{2}$. Note that it might be more profitable for a deviating firm to play the corner solution and earn $\pi^{dev} = \alpha\bar{p}$ than to unshroud and undercut the collusive price marginally. But since this would only decrease the collusion-to-deviation profit ratio and result in a higher critical discount factor, collusion would be further destabilized. We can therefore conclude, that coordinating on other than monopoly profits cannot stabilize collusion. $\square$

A.2. Appendix to Chapter II

A.2.1. Appendix

	Sessions	Groups	Subjects	Female	Beliefs	Gender revealed
Treatment						
Baseline	4	12	84	51 (60.7%)	no	no
Belief	4	12	84	47 (56.0%)	yes	no
Cheating	4	12	84	51 (60.7%)	no	no
Gender	4	12	84	43 (51.2%)	yes	yes
Sum	16	48	336	192 (57.1%)		

Table A.1.: Overview treatments

	(1)	(2)	(3)	(4)
Female	−3.436	−3.436	−4.501	−5.256
	(7.012)	(7.018)	(5.391)	(3.332)
Period		6.046^{***}	6.046^{***}	5.946^{***}
		(0.482)	(0.482)	(1.054)
Risk attitude			−1.194	−1.194
			(1.789)	(1.791)
Female x period				0.168
				(1.142)
Constant	114.9^{***}	87.66^{***}	94.00^{***}	94.45^{***}
	(6.595)	(5.241)	(7.532)	(9.877)
Observations	576	576	576	576
R^2	0.002	0.162	0.170	0.170

Dependent variable is achieved points. Robust standard errors clustered on session in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

Table A.2.: Random Effects GLS Regressions for the baseline treatment with achieved points as the dependent variable

	Random Effects GLS		Random Effects Tobit	
	(1)	(2)	(3)	(4)
Female	-14.62^{**}	37.64^{**}	-19.59^{***}	76.31^{*}
	(5.822)	(15.00)	(5.806)	(41.85)
Power	1.326^{**}	2.731	4.023	7.291^{*}
	(0.538)	(1.979)	(3.066)	(4.132)
Achievement	0.821	-1.505	1.248	-2.848
	(2.766)	(5.300)	(2.987)	(3.776)
Benevolence	0.152	3.932	2.273	9.652^{*}
	(2.514)	(2.752)	(3.479)	(5.083)
Security	0.511	4.807	-0.540	6.461
	(2.015)	(3.764)	(2.587)	(4.432)
Period	0.117	0.117	-0.0619	-0.0636
	(0.262)	(0.263)	(0.268)	(0.268)
Female x power		1.131		0.460
		(5.786)		(6.331)
Female x achievement		4.360		7.874
		(4.606)		(5.621)
Female x benevolence		-9.168^{***}		-16.75^{**}
		(2.470)		(6.995)
Female x security		-7.345^{*}		-11.98^{**}
		(4.365)		(5.454)
Constant	15.27	-15.69	-3.265	-62.05^{*}
	(16.39)	(16.26)	(20.08)	(33.55)
Observations	536	536	536	536
# of left censored			153	153
# right censored			21	21
Log Likelihood			-1617.572	-1612.0609
R^2	0.158	0.230		

Dependent variable is sabotage. Standard errors (for GLS Robust standard errors clustered on session) in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

Table A.3.: Random Effects GLS and Random Effects Tobit regressions for the baseline treatment with sabotage as the dependent variable

	Random Effects GLS			Random Effects Tobit		
	(1)	(2)	(3)	(4)	(5)	(6)
Dummy female	-10.52^{**}	-9.060^{*}	-5.800	-15.21^{***}	-13.01^{**}	-6.979
	(4.442)	(5.016)	(4.277)	(5.405)	(5.603)	(6.317)
Period		-0.268	0.175		-0.778^{**}	0.0193
		(0.452)	(0.383)		(0.328)	(0.504)
Risk attitude		1.127^{*}	1.127^{*}		1.692	1.692
		(0.681)	(0.681)		(1.277)	(1.279)
Female x period			-0.724^{**}			-1.368^{**}
			(0.323)			(0.662)
Constant	25.38^{***}	20.11^{***}	18.12^{***}	23.17^{***}	16.86^{*}	13.31
	(3.905)	(7.002)	(6.431)	(4.207)	(8.604)	(8.784)
Observations	576	576	576	576	576	576
# of left censored						186
# right censored						19
Log Likelihood				-1787.964	-1784.2817	-1782.1486
R^2	0.073	0.088	0.090			

Dependent variable is cheating. Standard errors (for GLS Robust standard errors clustered on session) in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

Table A.4.: Random Effects GLS and Random Effects Tobit Regressions for the cheating treatment with cheating as the dependent variable

	Random Effects GLS			Random Effects Tobit		
	(1)	(2)	(3)	(4)	(5)	(6)
Female	-12.96^{***}	-12.96^{***}	-9.797^{***}	-15.79^{***}	-15.51^{***}	-11.06^{***}
	(1.643)	(1.267)	(1.302)	(4.361)	(4.384)	(3.905)
Period		0.417	-0.339		0.359	-0.655^{**}
		(0.255)	(0.349)		(0.240)	(0.305)
Risk attitude		0.00385	0.0749		0.524	0.649
		(0.707)	(0.508)		(1.038)	(0.798)
Female x period			0.481			0.576
			(0.343)			(0.426)
Belief achieved			-0.00705			-0.0117
points			(0.0126)			(0.00976)
Belief sabotage			0.481^{***}			0.619^{***}
			(0.0587)			(0.0513)
Constant	25.56^{***}	23.67^{***}	14.85^{***}	23.44^{***}	19.01^{***}	8.095
	(1.577)	(3.144)	(3.163)	(3.111)	(6.497)	(5.221)
Observations	576	576	576	576	576	576
# of left censored				139	139	139
# right censored				4	4	4
Log Likelihood				-1878.4856	-1877.243	-1809.6775
R^2	0.133	0.136	0.461			

Dependent variable is sabotage. Standard errors (for GLS Robust standard errors clustered on session) in parentheses. ***p$<$0.01, **p$<$ 0.05, *p$<$ 0.1

Table A.5.: Random Effects GLS and Random Effects Tobit Regressions for the belief treatment with sabotage as the dependent variable

	Random Effects GLS		
	(1)	(2)	(3)
Female	-15.88^{***}	-15.81^{***}	-15.13^{***}
	(3.993)	(4.005)	(3.365)
Opponent female	4.854^{*}	4.390	6.517
	(2.891)	(3.024)	(4.403)
Female x opp. female			-5.070
			(3.950)
Period		6.574^{***}	6.370^{***}
		(0.383)	(0.212)
Female x period			0.445
			(0.291)
Constant	119.1^{***}	89.67^{***}	89.65^{***}
	(2.598)	(1.506)	(2.461)
Observations	568	568	568
R^2	0.073	0.299	0.300

Dependent variable is achieved points. Robust standard errors clustered on session in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

Table A.6.: Random Effects GLS and Random Effects Tobit Regressions in the gender treatment with achieved points as the dependent variable

	Random Effects GLS			Random Effects Tobit		
	(1)	(2)	(3)	(4)	(5)	(6)
Female	-9.859^{**}	-8.268^{*}	-5.550^{*}	-17.35^{***}	-14.47^{***}	-10.44^{**}
	(4.011)	(4.837)	(3.093)	(5.230)	(5.457)	(4.130)
Opponent female	-0.278	1.030	1.611	-0.390	1.677	2.465
	(0.624)	(1.356)	(1.539)	(1.438)	(1.810)	(1.712)
Female x opp. female		-3.013	-1.234		-5.411^{*}	-3.266
		(2.456)	(2.427)		(2.951)	(2.777)
Period		0.357	0.00985		0.350	-0.0748
		(0.562)	(0.477)		(0.255)	(0.310)
Belief sabotage			0.520^{***}			0.605^{***}
			(0.0421)			(0.0706)
Belief achieved			-0.00561			-0.0140
points			(0.00543)			(0.0240)
Constant	20.82^{***}	18.63^{***}	9.709^{***}	17.75^{***}	15.27^{***}	6.039^{*}
	(3.089)	(5.659)	(3.220)	(3.698)	(3.918)	(3.602)
Observations	568	568	568	568	568	568
# of left censored				195	195	195
# right censored				3	3	3
Log Likelihood				-1629.4484	-1626.9575	-1592.3773
R^2	0.091	0.093	0.460			

Dependent variable is sabotage. Standard errors (for GLS robust standard errors clustered on session) in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

Table A.7.: Random Effects GLS and Random Effects Tobit Regressions for the gender treatment with sabotage as the dependent variable

	Random Effects GLS				Random Effects Tobit	
	Achieved points		Sabotage			
	(1)	(2)	(3)	(4)	(5)	(6)
Female	-4.696	-4.868	-11.51^{***}	-11.33^{***}	-16.12^{***}	-15.88^{***}
	(3.531)	(3.325)	(1.769)	(1.694)	(2.523)	(2.539)
Belief treatment		-2.767		0.528		0.376
		(5.922)		(1.951)		(3.482)
Cheating treatment		-5.713		-0.709		-1.222
		(3.878)		(1.348)		(3.490)
Gender treatment		0.289		-3.937^{**}		-7.532^{**}
		(3.795)		(1.848)		(3.534)
Period		6.105^{***}		0.177		0.00990
		(0.237)		(0.189)		(0.135)
Risk attitude		-0.110		0.561		0.937^{*}
		(0.775)		(0.368)		(0.556)
Constant	113.6^{***}	88.77^{***}	24.43^{***}	21.77^{***}	22.17^{***}	19.41^{***}
	(3.006)	(5.116)	(1.515)	(2.553)	(1.871)	(4.131)
Observations	2,296	2,296	2,296	2,296	2,296	2,296
# of left censored					680	680
# right censored					47	47
Log Likelihood					-7069.3929	-7064.7676
R^2	0.005	0.178	0.102	0.117		

Dependent variable is achieved points / sabotage. Standard errors (for GLS: robust standard errors clustered on session) in parentheses. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$

Table A.8.: Random Effects GLS and Random Effects Tobit Regressions (only sabotage) with achieved points / sabotage as the dependent variable for all treatments

A.2.2. Instructions of the experiment (baseline treatment)

Welcome to this experiment!

You are participating in an economic experiment. All decisions are made privately, meaning that none of the other participants learns the identity of someone having made a decision. The payment is private information as well; none of the participants learns how much others have earned. Please read the instructions carefully. If you have trouble understanding the instructions, please take a second look at them. If you still have questions, please give us a signal.

Overview

- The experiment consists of 8 identical periods. At the end, one period will be drawn randomly for payment. After the last period, you will receive an overview of your possible payments.
- In this experiment, you are randomly assigned to a unit with seven members each. Throughout the experiment, you will only play with members from your unit. You and the other participants never learn the identity of the other members.
- There are two types of players: type I and type II. There are six players of type I and one player of type II in each unit.
- You will learn about your type at the beginning of the experiment. Your type will not change throughout the whole experiment.
- In each period, every participant encodes words into numbers. You have to replace each letter of a word with the numbers given in Table 1. You will earn one point for each letter.
- Before the actual 8 periods of the experiment begin, you have the opportunity to become familiar with the task in a one minute practice period. The practice period only serves better understanding of the task and does not affect your payment.
- Within a unit, two players of type I are assigned to one group. Please note that the group members change every period and that the identity of the player remains unknown.
- Within a group, the overall score of both type I players will be compared at the end of each period. The player with the higher overall score earns 500 tokens, the other player earns 200 tokens.
- Players of type II do not make any decisions that affect their own payment or the payment of other players in this experiment.

- At the end of the experiment, you will complete a short questionnaire. When all participants have completed the questionnaire, we will start paying the participants one by one.

Course of action

- Your task is to encode words into numbers. You have to replace each letter of a word with the numbers given in Table 1:
  - Example: The word *hat* appears on the screen. According to Table 1, $H = 15$, $A = 21$ and $T = 91$, so the code for *hat* is: 152191.
  - For each letter, you have to enter the assigned number in a separate box. You can go from box to box using the tab key.
- Five-digit, six-digit,and seven-digit words will appear. You will earn one point for each letter. Please note that you will only earn points, if you encode the whole word correctly. The sum of the points is the obtained score.
- When you have entered the code and pressed *OK*, you receive a notification whether the word was encoded correctly. If so, please click on *next* in order to receive a new word. If the word was not encoded correctly, please try again until you succeed. You have five minutes working time per period. Thereafter, working time will stop automatically.

Type I:

- In each period, you will be randomly assigned to some other type I player from your unit in a group.
- The other player also encodes words into numbers. Note that you will both receive the same words in the same order.
- At the end of each period, the overall score of both type I players will be compared. The player with the higher overall score receives the higher payment of 500 tokens. The more words a player has encoded correctly, the higher the obtained score will be. Please note that your overall score is only used for comparison with the score of the other player. Only if your overall score is higher than the score of the other player, you receive the higher payment of 500 tokens. It does not matter by how many points you outperform the other player. If your overall score is lower than the score of the other player, you receive 200 tokens.

- Before a period starts, you have the opportunity to reduce the overall score of the other type I player by the amount of X. In this way, the other type I player has a disadvantage when the overall scores are compared at the end of the period. The more points you deduct from the other player, the higher are your resulting costs. The costs will be deducted from your payment for this period in any case. An overview of the costs can be found in Table 2.

- In the same way, the other type I player decides whether he wants to reduce your score before the period starts.

- The overall score of a type I player consists of his obtained encoding score minus the amount of X the other player has chosen. Please note that the other type I player never learns which amount of X you have chosen, nor do you receive any information on the amount of X the other type I player has chosen.

  Overall score = obtained encoding score - amount of X (chosen by the other type I player)

- At the end of each period, we will show you the following information:
  - Your obtained encoding score
  - The amount of X you have chosen
  - Your payment, if this period is drawn.

- Thereafter a new period starts.

Type II:

- Type II players have no influence on the overall score of type I players. You will also encode words into numbers, but you do not receive a special payment for this. At the end of each period, you receive an overview of the overall score of all type I players from your unit. Your payment depends on the average overall score of all six type I players from your unit. One point equals two tokens.

Payment

- At the end of the experiment, the period that determines your payment is drawn randomly.

- During the whole experiment, the payments are shown in the currency *tokens*, which will be converted at the end. The conversion rate is 25 tokens $\rightarrow$ 1 Euro.

Type I:

- The overall score of both type I players from the allotted period influence their payment and the payment of type II. The type I player from the group with the higher overall score receives 500 tokens, the type I player with the lower overall score receives 200 tokens.

- The costs for the chosen amount of X in the allotted period will be deducted from the payment of each player. This yields the overall payment at the end of the experiment. In addition, each player receives a fix amount of 150 tokens.
  Higher overall score: **500 tokens - costs for the score of X + 150 tokens**
  Lower overall score: **200 tokens - costs for the score of X + 150 tokens**

Type II

- The type II player receives the average of the obtained overall score of the six type I players from his unit in the allotted period as a payment. One point equals 2 tokens.
  2x(average of the unit's overall score) + 100 tokens

Letter	Number	Letter	Number
A	21	N	32
B	54	O	56
C	13	P	10
D	67	Q	23
E	85	R	49
F	31	S	82
G	46	T	91
H	15	U	37
I	98	V	43
J	75	W	52
K	42	X	87
L	27	Y	93
M	19	Z	30

Table A.9.: Overview of the numercical codes used

There will be no costs if you choose X=0.										
X	1	2	3	4	5	6	7	8	9	10
Costs	0.07	0.29	0.64	1.14	1.79	2.57	3.50	4.57	5.79	7.14
X	11	12	13	14	15	16	17	18	19	20
Costs	8.64	10.29	12.07	14.00	16.07	18.29	20.64	23.14	25.79	28.57
X	21	22	23	24	25	26	27	28	29	30
Costs	31.50	34.57	37.79	41.14	44.64	48.29	52.07	56.00	60.07	64.29
X	31	32	33	34	35	36	37	38	39	40
Costs	68.64	73.14	77.79	82.57	87.50	92.57	97.79	103.14	108.64	114.29
X	41	42	43	44	45	46	47	48	49	50
Costs	120.07	126.00	132.07	138.29	144.64	151.14	157.79	164.57	171.50	178.57
X	51	52	53	54	55	56	57	58	59	60
Costs	185.79	193.14	200.64	208.29	216.07	224.00	232.07	240.29	248.64	257.14
X	61	62	63	64	65	66	67	68	69	70
Costs	265.79	274.57	283.50	292.57	301.79	311.14	320.64	330.29	340.07	350.00

Table A.10.: Overview of the costs for the score X

A.3. Appendix to Chapter III

Proof of Lemma 3.1. Suppose $\mathcal{U}$ is an outcome space with $N \geq 2$ elements, $u_1 > u_2 > \ldots > u_N$. Furthermore, let $L^A = \{p_1^A, p_2^A, \ldots, p_N^A\}$ and $L^B = \{p_1^B, p_2^B, \ldots, p_N^B\}$ denote two probability distributions over the set $\mathcal{U}$, where p_k^j denotes the probability that outcome u_k is realized under probability distribution L^j with $j \in \{A, B\}$. For $k = 1, \ldots, N$, $\mathcal{U}(k) = \{u_1, \ldots, u_k\}$ denotes the "truncated" outcome space which contains only the k highest elements of $\mathcal{U}$. For $k = 1, \ldots, N$ and $j \in \{A, B\}$, $L^j(k) = (p_1^j(k), \ldots, p_k^j(k))$ denotes the probability distribution over the truncated outcome space $\mathcal{U}(k)$ with $p_n^j(k) = p_n^j$ for $n < k$ and $p_k^j(k) = \sum_{n=k}^N p_n^j$. Differentiation of

$$f(L^A(k), L^B(k)) = \frac{1}{2}\sum_{s=1}^{k}\sum_{t=1}^{k} p_s^A(k)p_t^A(k)|u_s - u_t| \\ - \sum_{s=1}^{k}\sum_{t=1}^{k} p_s^A(k)p_t^B(k)|u_s - u_t| + \frac{1}{2}\sum_{s=1}^{k}\sum_{t=1}^{k} p_s^B(k)p_t^B(k)|u_s - u_t|$$

with respect to u_k yields

$$\frac{df(L^A(k), L^B(k))}{du_k} = [p_k^A(k) - p_k^B(k)]^2 \geq 0. \tag{A.3}$$

For $k = 1$, $\mathcal{U}(1) = \{u_1\}$ and the lotteries $L^A(1)$ and $L^B(1)$ are degenerate with $p_1^A(1) = p_1^B(1) = 1$. In consequence,

$$f(L^A(1), L^B(1)) = 0.$$

By (A.3),

$$f(L^A(k), L^B(k)) \leq f(L^A(k-1), L^B(k-1)), \qquad \forall k \geq 2, \tag{A.4}$$

where (A.4) holds with equality if and only if $p_k^A(k) = p_k^B(k)$. Hence, for $f(L^A, L^B) = f(L^A(N), L^B(N)) = 0$ to hold, we must have $p_k^A(k) = p_k^B(k)$ for all $k = 1, \ldots, N$. Given that $p_t^A(t) = p_t^B(t)$ for all $t = k+1, \ldots, N$, then $p_k^A(k) = p_k^B(k)$ if and only if $p_k^A = p_k^B$. Therefore $f(L^A, L^B) = 0$ holds if and only if L^A and L^B are identical, i.e., $p_k^A = p_k^B$ for all $k = 1, \ldots, N$. Conversely, if $p_k^A \neq p_k^B$ for some $k = 1, \ldots, N$, then $f(L^A, L^B) < 0$. □

Proof of Proposition 3.1. With $U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$ being a linear function of the components of σ^i, the derivative of $U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$ with respect to $\sigma^i(s_k^i)$ is linear in the components of $\hat{\sigma}^i$:

$$MU_k^i(\hat{\sigma}^i, \sigma^{-i}) := \frac{\partial U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})}{\partial \sigma^i(s_k^i)} = q_{k1}^i(\sigma^{-i})\hat{\sigma}^i(s_1^i) + ... + q_{kM^i}^i(\sigma^{-i})\hat{\sigma}^i(s_{M^i}^i) + b_k^i(\sigma^{-i})$$

with

$$q^i_{km}(\sigma^{-i}) =$$

$$\sum_{\theta\in\Theta} Q(\theta)\Big\{ \sum_{s^{-i}\in\mathcal{S}^{-i}} \big(\prod_{(s^j)_{j\neq i}\in\mathcal{S}^{-i}} \sigma^j(s^j)\big)\Big[\sum_{\tilde{\theta}\in\Theta} Q(\tilde{\theta})\Big(\sum_{\tilde{s}^{-i}\in\mathcal{S}^{-i}} \big(\prod_{j\neq i} \sigma^j(\tilde{s}^j)\big)\mu(u((s^i_k, s^{-i}), \theta) - u((s^i_m, \tilde{s}^{-i}), \tilde{\theta}))\Big)\Big]\Big\}$$

and

$$b^i_k(\sigma^{-i}) = \sum_{\theta\in\Theta} Q(\theta)\Big[\sum_{s^{-i}\in\mathcal{S}^{-i}} \big(\prod_{j\neq i} \sigma^j(s^j)\big) u((s^i_k, s^{-i}), \theta)\Big],$$

where $\mathcal{S}^{-i} = \times_{j\neq i}\mathcal{S}^j$. The coefficients $q^i_{km}(\sigma^{-i})$ as well as $b^i_k(\sigma^{-i})$ are continuous functions of the components of player i's opponents' strategies. Defining $a^i_{km}(\sigma^{-i}) = q^i_{km}(\sigma^{-i}) - b^i_k(\sigma^{-i})$, we can rewrite the system of M^i linear equations that represent player i's marginal utilities in matrix notation as follows:

$$\begin{pmatrix} MU^i_1(\hat{\sigma}^i, \sigma^{-i}) \\ \vdots \\ MU^i_{M^i}(\hat{\sigma}^i, \sigma^{-i}) \end{pmatrix} = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \dots & a_{1M^i}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{M^i1}(\sigma^{-i}) & \dots & a_{M^iM^i}(\sigma^{-i}) \end{pmatrix}}_{=:A(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s^i_1) \\ \vdots \\ \hat{\sigma}^i(s^i_{M^i}) \end{pmatrix}, \tag{A.5}$$

where the matrix $A(\sigma^{-i})$ depends only on the strategies of player i's opponents.

Suppose a mixed strategy $\bar{\sigma}^i$ with $|\Gamma(\bar{\sigma}^i)| = m \geq 2$ is a PE for player i given her opponents' strategy profile σ^{-i}. W.l.o.g., assume that $\bar{\sigma}^i$ assigns strictly positive probability to the first m pure strategies in $\mathcal{S}^i$, i.e., $\bar{\sigma}^i(s^i_k) > 0$ for $k = 1, \dots, m$ and $\bar{\sigma}^i(s^i_k) = 0$ for $k > m$, where $\sum_{k=1}^m \bar{\sigma}^i(s^i_k) = 1$. As described in the text, $MU^i_k(\hat{\sigma}^i, \sigma^{-i})$ reflects the attractiveness to play pure strategy s^i_k. Since $\bar{\sigma}^i$ is assumed to be a PE, $MU^i_1(\bar{\sigma}^i, \sigma^{-i}) = \dots = MU^i_m(\bar{\sigma}^i, \sigma^{-i}) = \bar{u} \geq \max_{k>m} MU^i_k(\bar{\sigma}^i, \sigma^{-i})$. With $\bar{\sigma}^i(s^i_k) = 0$ for $k > m$, the mixing probabilities $\bar{\sigma}^i(s^i_1), \dots, \bar{\sigma}^i(s^i_m)$ are thus a solution of the following system of linear equations:

$$\begin{pmatrix} \bar{u} \\ \vdots \\ \bar{u} \end{pmatrix} = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \dots & a_{1m}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{m1}(\sigma^{-i}) & \dots & a_{mm}(\sigma^{-i}) \end{pmatrix}}_{=:A'(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s^i_1) \\ \vdots \\ \hat{\sigma}^i(s^i_m) \end{pmatrix}. \tag{A.6}$$

Based on these observations, we will prove the two parts of the statement in turn.

(i) The proof proceeds in two steps: first, we show that the statement holds if matrix $A'(\sigma^{-i})$ has full rank; second, we show that no pure strategy in $\Gamma(\bar{\sigma}^i)$ being redundant implies full rank of matrix $A'(\sigma^{-i})$.

STEP 1: Suppose matrix $A'(\sigma^{-i})$ has full rank. Then the system of linear equations in (A.6) has a unique solution, which (by hypothesis) is given by the vector $(\bar{\sigma}^i(s^i_1), \dots, \bar{\sigma}^i(s^i_m))$ with $\sum_{k=1}^m \sigma^i(s^i_k) = 1$. In contradiction to the statement, suppose that there exists a different PE,

$(\tilde{\sigma}^i(s_1^i), \ldots, \tilde{\sigma}^i(s_{M^i}^i)) \neq (\bar{\sigma}^i(s_1^i), \ldots, \bar{\sigma}^i(s_{M^i}^i))$, that mixes over the same set of pure strategies, i.e., $\tilde{\sigma}^i(s_k^i) > 0$ for $k = 1, \ldots, m$ and $\tilde{\sigma}^i(s_k^i) = 0$ for $k > m$, where $\sum_{k=1}^m \tilde{\sigma}^i(s_k^i) = 1$. By the logic described above, the vector $(\tilde{\sigma}^i(s_1^i), \ldots, \tilde{\sigma}^i(s_m^i))$ solves a system of linear equations

$$\begin{pmatrix} \tilde{u} \\ \vdots \\ \tilde{u} \end{pmatrix} = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \ldots & a_{1m}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{m1}(\sigma^{-i}) & \ldots & a_{mm}(\sigma^{-i}) \end{pmatrix}}_{=A'(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s_1^i) \\ \vdots \\ \hat{\sigma}^i(s_m^i) \end{pmatrix}. \tag{A.7}$$

By full rank of $A'(\sigma^{-i})$, we must have $\tilde{u} \neq \bar{u}$, because otherwise $(\tilde{\sigma}^i(s_1^i), \ldots, \tilde{\sigma}^i(s_m^i)) = (\bar{\sigma}^i(s_1^i), \ldots, \bar{\sigma}^i(s_m^i))$. Hence either $\bar{u}$ or $\tilde{u}$ differs from zero. Suppose, that $\tilde{u} \neq 0$. In consequence, (A.7) implies that $(\bar{\sigma}^i(s_1^i), \ldots, \bar{\sigma}^i(s_m^i)) = (\frac{\bar{u}}{\tilde{u}}\tilde{\sigma}^i(s_1^i), \ldots, \frac{\bar{u}}{\tilde{u}}\tilde{\sigma}^i(s_m^i))$. But then $1 = \sum_{k=1}^m \bar{\sigma}^i(s_k^i) = \frac{\bar{u}}{\tilde{u}} \sum_{k=1}^m \tilde{\sigma}^i(s_k^i)$ implies $\sum_{k=1}^m \tilde{\sigma}^i(s_k^i) \neq 1$—a contradiction. The same line of reasoning applies for $\bar{u} \neq 0$.

STEP 2: Suppose matrix $A'(\sigma^{-i})$ does not have full rank, i.e., one of the row vectors of matrix $A'(\sigma^{-i})$ is a linear combination of the other row vectors. Let the k-th row vector be denoted by $a_k(\sigma^{-i}) = (a_{k1}(\sigma^{-i}), \ldots, a_{km}(\sigma^{-i}))$. Suppose, w.l.o.g., that the last row vector $a_m(\sigma^{-i})$ can be expressed as a linear combination of the row vectors $a_1(\sigma^{-i}), \ldots, a_{m-1}(\sigma^{-i})$, i.e., $a_m(\sigma^{-i}) = \sum_{k=1}^{m-1} \gamma_k a_k(\sigma^{-i})$ for some numbers $\gamma_1, \ldots, \gamma_{m-1}$. Therefore

$$\begin{aligned} MU_m^i(\hat{\sigma}^i, \sigma^{-i}) &= a_{m1}(\sigma^{-i})\hat{\sigma}^i(s_1^i) + \ldots + a_{mm}(\sigma^{-i})\hat{\sigma}^i(s_m^i) \\ &= \left[\sum_{k=1}^{m-1} \gamma_k a_{k1}(\sigma^{-i})\right] \hat{\sigma}^i(s_1^i) + \ldots + \left[\sum_{k=1}^{m-1} \gamma_k a_{km}(\sigma^{-i})\right] \hat{\sigma}^i(s_m^i) \\ &= \sum_{k=1}^{m-1} \gamma_k \left[a_{k1}(\sigma^{-i})\hat{\sigma}^i(s_1^i) + \ldots + a_{km}(\sigma^{-i})\hat{\sigma}^i(s_m^i)\right] \\ &= \sum_{k=1}^{m-1} \gamma_k MU_k^i(\hat{\sigma}^i, \sigma^{-i}) \end{aligned}$$

holds for every $\hat{\sigma}^i$ with $\Gamma(\hat{\sigma}^i) \subseteq \Gamma(\bar{\sigma}^i)$. Since $MU_k^i(\bar{\sigma}^i, \sigma^{-i}) = \bar{u}$ for $k = 1, \ldots, m$, this immediately implies $\sum_{k=1}^{m-1} \gamma_k = 1$. Since marginal utilities of pure strategies are constant given $\hat{\sigma}^i$ and σ^{-i}, for any σ^i with $\Gamma(\sigma^i) \subseteq \Gamma(\bar{\sigma}^i)$ we thus have

$$\begin{aligned} U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i}) &= MU_1^i(\hat{\sigma}^i, \sigma^{-i})\sigma^i(s_1^i) + \ldots + MU_m^i(\hat{\sigma}^i, \sigma^{-i})\sigma^i(s_m^i) \\ &= \sum_{k=1}^{m-1} MU_k^i(\hat{\sigma}^i, \sigma^{-i}) \left[\sigma^i(s_k^i) + x\gamma_k \sigma^i(s_m^i)\right] + MU_m^i(\hat{\sigma}^i, \sigma^{-i})(1-x)\sigma^i(s_m^i), \end{aligned}$$

for all $x \in [0,1]$. Consider the mixed strategy $\sigma_{\tilde{x}}^i = (\sigma_{\tilde{x}}^i(s_1^i), \ldots, \sigma_{\tilde{x}}^i(s_{M^i}^i))$ with

$$\sigma_{\tilde{x}}^i(s_k^i) = \begin{cases} \bar{\sigma}^i(s_k^i) + \tilde{x}\gamma_k \bar{\sigma}^i(s_m^i) & \text{if} \quad k \leq m-1 \\ (1-\tilde{x})\bar{\sigma}^i(s_m^i) & \text{if} \quad k = m \\ 0 & \text{if} \quad k > m \end{cases},$$

where

$$\tilde{x} = \min\left\{1\,,\ \min_{k\in\{\tilde{k}\,|\,1\leq\tilde{k}\leq m-1\,,\,\gamma_{\tilde{k}}<0\}}\left\{-\frac{\sigma^i(s_k^i)}{\gamma_k\sigma^i(s_m^i)}\right\}\right\}.$$

Note that $\sum_{k=1}^{m-1}\gamma_k = 1$ implies $\sum_{k=1}^{M^i}\sigma_{\tilde{x}}^i(s_k^i) = 1$. By choice of $\tilde{x}$, we also have $\sigma_{\tilde{x}}^i(s_k^i) \geq 0$ for all $k = 1,\ldots,m$ and $\sigma_{\tilde{x}}^i(s_k^i) = 0$ for at least one $k = 1,\ldots,m$. Overall, strategy $\sigma_{\tilde{x}}^i$ yields utility $U^i(\sigma_{\tilde{x}}^i,\hat{\sigma}^i,\sigma^{-i}) = U^i(\bar{\sigma}^i,\hat{\sigma}^i,\sigma^{-i})$ for all $\hat{\sigma}^i$ with $\Gamma(\hat{\sigma}^i)\subseteq\Gamma(\bar{\sigma}^i)$. With

$$GL^i(\sigma^i,\hat{\sigma}^i,\sigma^{-i}) \equiv \sum_{u\in\mathcal{U}^i}\sum_{\tilde{u}\in\mathcal{U}^i}P^i(u|(\sigma^i,\sigma^{-i}))\cdot P^i(\tilde{u}|(\hat{\sigma}^i,\sigma^{-i}))\cdot\mu(u-\tilde{u}),$$

we obtain that $U^i(\sigma_{\tilde{x}}^i,\sigma_{\tilde{x}}^i,\sigma^{-i}) = U^i(\bar{\sigma}^i,\sigma_{\tilde{x}}^i,\sigma^{-i})$ if and only if

$$E[L^i(\bar{\sigma}^i,\sigma^{-i})] - E[L^i(\sigma_{\tilde{x}}^i,\sigma^{-i})] \;=\; GL^i(\sigma_{\tilde{x}}^i,\sigma_{\tilde{x}}^i,\sigma^{-i}) - GL^i(\bar{\sigma}^i,\sigma_{\tilde{x}}^i,\sigma^{-i}). \tag{A.8}$$

Likewise, $U^i(\bar{\sigma}^i,\bar{\sigma}^i,\sigma^{-i}) = U^i(\sigma_{\tilde{x}}^i,\bar{\sigma}^i,\sigma^{-i})$ if and only if

$$E[L^i(\bar{\sigma}^i,\sigma^{-i})] - E[L^i(\sigma_{\tilde{x}}^i,\sigma^{-i})] \;=\; GL^i(\sigma_{\tilde{x}}^i,\bar{\sigma}^i,\sigma^{-i}) - GL^i(\bar{\sigma}^i,\bar{\sigma}^i,\sigma^{-i}). \tag{A.9}$$

(A.8) and (A.9) together imply

$$\begin{aligned}
&GL^i(\sigma_{\tilde{x}}^i,\bar{\sigma}^i,\sigma^{-i}) - GL^i(\bar{\sigma}^i,\bar{\sigma}^i,\sigma^{-i}) - GL^i(\sigma_{\tilde{x}}^i,\sigma_{\tilde{x}}^i,\sigma^{-i}) + GL^i(\bar{\sigma}^i,\sigma_{\tilde{x}}^i,\sigma^{-i}) = 0 \qquad \text{(A.10)}\\
&\Leftrightarrow \frac{1}{2}\sum_{u\in\mathcal{U}^i}\sum_{\tilde{u}\in\mathcal{U}^i}P^i(u|(\bar{\sigma}^i,\sigma^{-i}))\cdot P^i(\tilde{u}|(\bar{\sigma}^i,\sigma^{-i}))\cdot|u-\tilde{u}|\\
&-\sum_{u\in\mathcal{U}^i}\sum_{\tilde{u}\in\mathcal{U}^i}P^i(u|(\bar{\sigma}^i,\sigma^{-i}))\cdot P^i(\tilde{u}|(\sigma_{\tilde{x}}^i,\sigma^{-i}))\cdot|u-\tilde{u}|\\
&+\frac{1}{2}\sum_{u\in\mathcal{U}^i}\sum_{\tilde{u}\in\mathcal{U}^i}P^i(u|(\sigma_{\tilde{x}}^i,\sigma^{-i}))\cdot P^i(\tilde{u}|(\sigma_{\tilde{x}}^i,\sigma^{-i}))\cdot|u-\tilde{u}| = 0.
\end{aligned}$$

By Lemma 3.1, this holds if and only if $L^i(\bar{\sigma}^i,\sigma^{-i})$ and $L^i(\sigma_{\tilde{x}}^i,\sigma^{-i})$ are identical. Let w.l.o.g. the strategy being played with positive probability under $\bar{\sigma}^i$ and with zero probability under $\sigma_{\tilde{x}}^i$ be s_m^i. From $L^i(\bar{\sigma}^i,\sigma^{-i})$ and $L^i(\sigma_{\tilde{x}}^i,\sigma^{-i})$ being identical it follows that

$$\begin{aligned}
\sum_{j=1}^{m}\bar{\sigma}^i(s_j^i)L^i(s_j^i,\sigma^{-i}) &= \sum_{j=1}^{m-1}\sigma_{\tilde{x}}^i(s_j^i)L^i(s_j^i,\sigma^{-i})\\
\Leftrightarrow L^i(s_m^i,\sigma^{-i}) &= \sum_{j=1}^{m-1}\frac{\sigma_{\tilde{x}}^i(s_j^i)-\bar{\sigma}^i(s_j^i)}{\bar{\sigma}^i(s_m^i)}L^i(s_j^i,\sigma^{-i}).
\end{aligned}$$

The lottery that is induced by the pure strategy being played with zero probability under $\sigma_{\tilde{x}}^i$ and with positive probability under $\bar{\sigma}^i$ is a linear combination of the lotteries that are induced by the other pure strategies being played with positive probability with $\gamma(s_j^i) = \frac{\sigma_{\tilde{x}}^i(s_j^i)-\bar{\sigma}^i(s_j^i)}{\bar{\sigma}^i(s_m^i)}$ for $j = 1,\ldots,m-1$ and $\gamma(s_j^i) = 0$ for $j = m$. Hence, pure strategy s_m^i is redundant.

(ii) By Step 2 of part (i) of this proof, we can conclude that $A'(\bar{\sigma}^{-i})$ has full rank. The function $z(\hat{\sigma}^i, \sigma^{-i})$, defined by

$$z(\hat{\sigma}^i, \sigma^{-i}) = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \dots & a_{1m}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{m1}(\sigma^{-i}) & \dots & a_{mm}(\sigma^{-i}) \end{pmatrix}}_{=A'(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s_1^i) \\ \vdots \\ \hat{\sigma}^i(s_m^i) \end{pmatrix} - \begin{pmatrix} \bar{u} \\ \vdots \\ \bar{u} \end{pmatrix}, \tag{A.11}$$

is a C^1 function and its Jacobian with respect to the first m components of $\hat{\sigma}^i$ is invertible in an environment of its zero $\bar{\sigma}^i$. As a consequence of the implicit function theorem there exists a C^1 function $g : \Sigma^{-i} \to \mathbb{R}^m$ such that $z(g(\sigma^{-i}), \sigma^{-i}) = 0$ in an environment of $\bar{\sigma}^i$. Consider any σ_ε^{-i} such that $||\sigma_\varepsilon^{-i} - \sigma^{-i}|| \leq \varepsilon$ for some small $\varepsilon > 0$. By hypothesis, $MU_1^i(\bar{\sigma}^i, \sigma^{-i}) = \ldots = MU_m^i(\bar{\sigma}^i, \sigma^{-i}) > \max_{k>m} MU_k^i(\bar{\sigma}^i, \sigma^{-i})$, $\bar{\sigma}^i(s_k^i) > 0$ for $k \leq m$, and $\bar{\sigma}^i(s_k^i) = 0$ for $k > m$. Then the components of $g(\sigma_\varepsilon^{-i})$, which we denote by $(\hat{\sigma}_\varepsilon^i(s_1^i), \ldots, \hat{\sigma}_\varepsilon^i(s_m^i))$, are also strictly positive. Hence, for the mixed strategy $\sigma_\varepsilon^i = (\sigma_\varepsilon^i(s_1^i), \ldots, \sigma_\varepsilon^i(s_{M^i}^i))$ with

$$\sigma_\varepsilon^i(s_k^i) = \begin{cases} \frac{\hat{\sigma}_\varepsilon^i(s_k^i)}{\sum_{j=1}^m \hat{\sigma}_\varepsilon^i(s_j^i)} & \text{if} \quad k \leq m \\ 0 & \text{if} \quad k > m \end{cases}, \tag{A.12}$$

$A(\sigma_\varepsilon^{-i})\sigma_\varepsilon^i$ yields a vector of marginal utilities with $MU_1^i(\sigma_\varepsilon^i, \sigma_\varepsilon^{-i}) = \ldots = MU_m^i(\sigma_\varepsilon^i, \sigma_\varepsilon^{-i}) = \frac{\bar{u}}{\sum_{j=1}^m \hat{\sigma}_\varepsilon^i(s_j^i)} > \max_{k>m} MU_k^i(\sigma_\varepsilon^i, \sigma_\varepsilon^{-i})$. Thus, $\sigma_\varepsilon^i \in R(\sigma_\varepsilon^{-i})$ with $|\Gamma(\sigma_\varepsilon^i)| = m$. $\square$

Proof of Proposition 3.2. Suppose $\bar{\sigma}^i \in \Sigma^i$ is a mixed CPE with $\bar{\sigma}^i(s_k^i) > 0$ for $k = 1, \ldots, m$ (where $m \geq 2$) and $\bar{\sigma}^i(s_k^i) = 0$ for $k > m$. (Assuming that player i mixes over the first m pure strategies is without loss of generality, because we can always relabel strategies.) Furthermore, for $1 \leq m', m'' \leq m$ and $m' \neq m''$, let the two strategies $\sigma_{m'}$ and $\sigma_{m''}$ be defined by

$$\sigma_{m'}(s_{m'}^i) = \bar{\sigma}^i(s_{m'}^i) + \bar{\sigma}^i(s_{m''}^i),\ \sigma_{m'}(s_{m''}^i) = 0,\ \sigma_{m'}(s_k^i) = \bar{\sigma}^i(s_k^i) \text{ for } k \neq m', m'' \tag{A.13}$$

and

$$\sigma_{m''}(s_{m'}^i) = 0,\ \sigma_{m''}(s_{m''}^i) = \bar{\sigma}^i(s_{m'}^i) + \bar{\sigma}^i(s_{m''}^i),\ \sigma_{m''}(s_k^i) = \bar{\sigma}^i(s_k^i) \text{ for } k \neq m', m'', \tag{A.14}$$

respectively. Thus, $\bar{\sigma}^i$ can be expressed as a convex combination of strategies $\sigma_{m'}$ and $\sigma_{m''}$,

$$\bar{\sigma}^i = \beta\sigma_{m'} + (1-\beta)\sigma_{m''},$$

where $\beta = \bar{\sigma}^i(s_{m'}^i)/[\bar{\sigma}^i(s_{m'}^i) + \bar{\sigma}^i(s_{m''}^i)]$. Since

$$U^i(\bar{\sigma}^i, \bar{\sigma}^i, \sigma^{-i}) = U^i(\beta\sigma_{m'} + (1-\beta)\sigma_{m''}, \beta\sigma_{m'} + (1-\beta)\sigma_{m''}, \sigma^{-i}),$$

player i strictly prefers playing either strategy $\sigma_{m'}$ or $\sigma_{m''}$ instead of playing $\bar{\sigma}^i$ if

$$\begin{aligned}
&\beta U^i(\sigma_{m'},\sigma_{m'},\sigma^{-i})+(1-\beta)U^i(\sigma_{m''},\sigma_{m''},\sigma^{-i})\\
&\quad> U^i(\beta\sigma_{m'}+(1-\beta)\sigma_{m''},\beta\sigma_{m'}+(1-\beta)\sigma_{m''},\sigma^{-i})\\
\Leftrightarrow&\beta GL^i(\sigma_{m'},\sigma_{m'},\sigma^{-i})+(1-\beta)GL^i(\sigma_{m''},\sigma_{m''},\sigma^{-i})\\
&\quad> GL^i(\beta\sigma_{m'}+(1-\beta)\sigma_{m''},\beta\sigma_{m'}+(1-\beta)\sigma_{m''},\sigma^{-i})\\
&\quad=\beta^2GL^i(\sigma_{m'},\sigma_{m'},\sigma^{-i})+(1-\beta)^2GL^i(\sigma_{m''},\sigma_{m''},\sigma^{-i})\\
&\quad+\beta(1-\beta)[GL^i(\sigma_{m'},\sigma_{m''},\sigma^{-i})+GL^i(\sigma_{m''},\sigma_{m'},\sigma^{-i})]\\
\Leftrightarrow&-GL^i(\sigma_{m'},\sigma_{m'},\sigma^{-i})+GL^i(\sigma_{m''},\sigma_{m'},\sigma^{-i})\\
&\quad+GL^i(\sigma_{m'},\sigma_{m''},\sigma^{-i})-GL^i(\sigma_{m''},\sigma_{m''},\sigma^{-i})<0 \qquad\qquad (A.15)\\
\Leftrightarrow&\frac{1}{2}\sum_{u\in\mathcal{U}^i}\sum_{\tilde{u}\in\mathcal{U}^i}P^i(u|(\sigma^i_{m'},\sigma^{-i}))\cdot P^i(\tilde{u}|(\sigma^i_{m'},\sigma^{-i}))\cdot|u-\tilde{u}|\\
&\quad-\sum_{u\in\mathcal{U}^i}\sum_{\tilde{u}\in\mathcal{U}^i}P^i(u|(\bar{\sigma}^i_{m'},\sigma^{-i}))\cdot P^i(\tilde{u}|(\sigma^i_{m''},\sigma^{-i}))\cdot|u-\tilde{u}|\\
&\quad+\frac{1}{2}\sum_{u\in\mathcal{U}^i}\sum_{\tilde{u}\in\mathcal{U}^i}P^i(u|(\sigma^i_{m''},\sigma^{-i}))\cdot P^i(\tilde{u}|(\sigma^i_{m''},\sigma^{-i}))\cdot|u-\tilde{u}|<0.
\end{aligned}$$

By Lemma 3.1, this last inequality holds if and only if $P^i(u|\sigma_{m'},\sigma^{-i})\neq P^i(u|\sigma_{m''},\sigma^{-i})$ for some $u\in\mathcal{U}^i$. With

$$P^i(u|\sigma_{m'},\sigma^{-i})=\sum_{k=1}^{M^i}\sigma_{m'}(s^i_k)P^i(u|s^i_k,\sigma^{-i})$$

and

$$P^i(u|\sigma_{m''},\sigma^{-i})=\sum_{k=1}^{M^i}\sigma_{m''}(s^i_k)P^i(u|s^i_k,\sigma^{-i}),$$

by (A.13) and (A.14) we have

$$\begin{aligned}
&P^i(u|\sigma_{m'},\sigma^{-i})\neq P^i(u|\sigma_{m''},\sigma^{-i})\\
&\Leftrightarrow\ \sigma_{m'}(s^i_{m'})P^i(u|s^i_{m'},\sigma^{-i})\neq\sigma_{m''}(s^i_{m''})P^i(u|s^i_{m''},\sigma^{-i})\\
&\Leftrightarrow\ P^i(u|s^i_{m'},\sigma^{-i})\neq P^i(u|s^i_{m''},\sigma^{-i}).
\end{aligned}$$

Hence, for $\bar{\sigma}^i$ to be a CPE it must hold that $P^i(u|s^i_{m'},\sigma^{-i})=P^i(u|s^i_{m''},\sigma^{-i})$ for any outcome $u\in\mathcal{U}^i$. Overall player i is only willing to mix between two actions if they induce the same lotteries over outcomes. □

Proof of Proposition 3.3. $(i)\Leftrightarrow(ii)$: Suppose the pure strategy profile (s^i,s^{-i}) constitutes a Nash equilibrium (NE). Player i's expected utility from playing and expecting to play some pure strategy s_k^i equals her material utility outcome under strategy profile (s_k^i, s^{-i}) and state of the world $\tilde{\theta}$, i.e., $U^i(s_k^i, s_k^i, s^{-i}) = u^i((s_k^i, s^{-i}), \tilde{\theta})$. By the definition of NE, $u^i((s^i, s^{-i}), \tilde{\theta}) \geq u^i((s_k^i, s^{-i}), \tilde{\theta})$ for all $k = 1, \ldots, M^i$, such that $U^i(s^i, s^i, s^{-i}) \geq U^i(s_k^i, s_k^i, s^{-i})$ for all $k = 1, \ldots, M^i$. Together with Proposition 3.2, i.e., the reluctance to play mixed strategies, this implies that the strategy played by player i in a given pure strategy NE also is a CPE for player i given the strategies of her opponents. Hence, any pure strategy NE is a CPNE by Definition 3.2.

Conversely, suppose the pure strategy profile (s^i, s^{-i}) constitutes a CPNE. By Definition 3.2, then $U^i(s^i, s^i, s^{-i}) \geq U^i(s_k^i, s_k^i, s^{-i})$ for all $k = 1, \ldots, M^i$, which implies that $u^i((s^i, s^{-i}), \tilde{\theta}) \geq u^i((s_k^i, s^{-i}), \tilde{\theta})$ for all $k = 1, \ldots, M^i$. Thus, the strategy played (and expected to be played) by player i in a pure strategy CPNE is a Nash best response given her opponents' strategies. Hence, any pure strategy CPNE is a NE.

$(i) \Leftrightarrow (iii)$: Suppose the pure strategy profile (s^i, s^{-i}) constitutes a NE. Given player i expects to play pure strategy s^i, deviating to any other pure strategy $s_k^i \neq s^i$ cannot be profitable for a loss-averse player. The reason is that she would incur not only (weakly) lower material utility—$u^i((s^i, s^{-i}), \tilde{\theta}) \geq u^i((s_k^i, s^{-i}), \tilde{\theta})$ for all $k = 1, ..., M^i$ by definition of NE—but also (weakly) lower psychological utility—she expected to obtain the maximum material utility $u^i((s^i, s^{-i}), \tilde{\theta})$ with certainty. By the same reasoning, deviating to a mixed strategy which involves some pure strategies that yield (weakly) lower material utility also is not profitable, because also psychological utility would be lower as any comparison with the deterministic reference point results in a (weak) loss. Thus, $U^i(s^i, s^i, s^{-i}) \geq U^i(\sigma^i, s^i, s^{-i})$ for all $\sigma^i \in \Sigma^i$, i.e., a strategy played by player i in a given pure strategy NE also is a PE for player i given the strategies of her opponents. Hence, any pure strategy NE is a PNE by Definition 3.1.

Conversely, suppose the pure strategy profile (s^i, s^{-i}) constitutes a PNE. Furthermore, in contradiction, suppose that (s^i, s^{-i}) does not constitute a NE. Then for some player, say player i, there must be some pure strategy s_k^i that yields strictly higher material utility than pure strategy s^i given s^{-i}, i.e., $u^i(s_k^i, s^{-i}, \tilde{\theta}) > u^i(s^i, s^{-i}, \tilde{\theta})$. A deviation to pure strategy s_k^i, however, represents a strictly profitable deviation for a loss-averse player i because it induces strictly higher material utility and also a strictly positive deterministic gain. This, however, contradicts the assumption that strategy profile (s^i, s^{-i}) constitutes a PNE. Hence, any pure strategy PNE is a NE. $\square$

Proof of Proposition 3.4. We prove both parts of the proposition in turn:

(i) We are going to show that $U^i(s^i, s^i, s^{-i}) \geq U^i(\sigma^i, s^i, s^{-i})$ for all $\sigma^i \in \Sigma^i$. To this end, note that

$$U^i(s^i,s^i,s^{-i})-U^i(\sigma^i,s^i,s^{-i})=\sum_{\tilde{s}^i\in\Gamma(\sigma^i)}\sigma^i(\tilde{s}^i)\Big\{\sum_{\theta\in\Theta}Q(\theta)\left[u^i((s^i,s^{-i}),\theta)-u^i((\tilde{s}^i,s^{-i}),\theta)\right]$$
$$+\sum_{\theta\in\Theta}Q(\theta)\sum_{\tilde{\theta}\in\Theta}Q(\tilde{\theta})\left[\mu\left(u^i((s^i,s^{-i}),\theta)-u^i((s^i,s^{-i}),\tilde{\theta})\right)-\mu\left(u^i((\tilde{s}^i,s^{-i}),\theta)-u^i((s^i,s^{-i}),\tilde{\theta})\right)\right]\Big\}. \tag{A.16}$$

With $u^i((s^i,s^{-i}),\theta)\geq u^i((\tilde{s}^i,s^{-i}),\theta)$ for all $\tilde{s}^i\in\mathcal{S}^i$, $s^{-i}\in\mathcal{S}^{-i}$, and $\theta\in\Theta$, it follows that $U^i(s^i,s^i,s^{-i})-U^i(\sigma^i,s^i,s^{-i})\geq 0$ for all $\sigma^i\in\Sigma^i$ by $\mu(\cdot)$ being strictly increasing.

For the reverse direction, it suffices to show that $U(s^i,\sigma^i,\sigma^{-i})>U(\sigma^i,\sigma^i,\sigma^{-i})$ for all $\sigma^i\in\Sigma^i\setminus\{s^i\}$. Irrespective of nature's draw and opponents' play the deviation to the (materially) weakly dominant strategy yields a weakly higher material utility. Hence, all losses are reduced or turned into gains and all gains are improved. Moreover, given any σ^{-i} there is a strict improvement in at least one gain or loss in material utility for at least one draw of nature. The (materially) weakly dominant strategy is, thus, strictly preferred in terms of expected material and psychological utility.

(ii) Denote by $L^1(s^1_1,\sigma^2)$ and $L^1(s^1_2,\sigma^2)$ the payoff lotteries for player 1 that are induced if he plays s^1_1 and s^1_2, respectively. Since $\sigma^1(s^1_2)=1-\sigma^1(s^1_1)$, the utility of player one of playing σ^1 when expecting to play $\hat{\sigma}^1$ is given by:

$$\begin{aligned}
&U^1(\sigma^1,\hat{\sigma}^1,\sigma^2)\\
&=\sigma^1(s^1_1)E[L^1(s^1_1,\sigma^2)]+(1-\sigma^1(s^1_1))E[L^1(s^1_2,\sigma^2)]+GL^1(\sigma^1,\hat{\sigma}^1,\sigma^2)\\
&=\sigma^1(s^1_1)E[L^1(s^1_1,\sigma^2)]+(1-\sigma^1(s^1_1))E[L^1(s^1_2,\sigma^2)]\\
&+\sigma^1(s^1_1)\hat{\sigma}^1(s^1_1)GL^1(s^1_1,s^1_1,\sigma^2)+(1-\sigma^1(s^1_1))\hat{\sigma}^1(s^1_1)GL^1(s^1_2,s^1_1,\sigma^2)\\
&+\sigma^1(s^1_1)(1-\hat{\sigma}^1(s^1_1))GL^1(s^1_1,s^1_2,\sigma^2)+(1-\sigma^1(s^1_1))(1-\hat{\sigma}^1(s^1_1))GL^1(s^1_2,s^1_2,\sigma^2)
\end{aligned}$$

Taking the derivative with respect to $\sigma(s^1_1)$ yields:

$$\begin{aligned}
&\frac{\partial U(\sigma^1,\hat{\sigma}^1,\sigma^2)}{\partial\sigma^1(s^1_1)}\\
&=E[L^1(s^1_1,\sigma^2)]-E[L^1(s^1_2,\sigma^2)]+\hat{\sigma}^1(s^1_1)GL^1(s^1_1,s^1_1,\sigma^2)-\hat{\sigma}^1(s^1_1)GL^1(s^1_2,s^1_1,\sigma^2)\\
&+(1-\hat{\sigma}^1(s^1_1))GL^1(s^1_1,s^1_2,\sigma^2)-(1-\hat{\sigma}^1(s^1_1))GL^1(s^1_2,s^1_2,\sigma^2)\\
&=E[L^1(s^1_1,\sigma^2)]-E[L^1(s^1_2,\sigma^2)]+GL^1(s^1_1,s^1_2,\sigma^2)-GL^1(s^1_2,s^1_2,\sigma^2) \qquad \text{(A.17)}\\
&+\hat{\sigma}^1(s^1_1)[GL^1(s^1_1,s^1_1,\sigma^2)-GL^1(s^1_2,s^1_1,\sigma^2)-GL^1(s^1_1,s^1_2,\sigma^2)+GL(s^1_2,s^1_2,\sigma^2)]
\end{aligned}$$

Suppose that s^1_1 and s^1_2 are not redundant for all $\sigma^2\in\Sigma^2$. By Lemma 3.1 the coefficient of $\hat{\sigma}^1(s^1_1)$ is then strictly positive. Whenever $\partial U(\sigma^1,\hat{\sigma}^1,\sigma^2)/\partial\sigma^1(s^1_1)=0$ for some $\hat{\sigma}^1(s^1_1)\in[0,1]$, player 1 is indifferent between all her mixed strategies given she expects to play $\hat{\sigma}^1(s^1_1)$. Hence, it is a PE for her to play $\sigma^1(s^1_1)=\hat{\sigma}^1(s^1_1)$. To characterize the complete set of PEs for player 1, define the function

$$h(\sigma^2(s_1^2)) = \frac{E[L^1(s_2^1,\sigma^2)] - E[L^1(s_1^1,\sigma^2)] - GL^1(s_1^1,s_2^1,\sigma^2) + GL^1(s_2^1,s_2^1,\sigma^2)}{GL^1(s_1^1,s_1^1,\sigma^2) - GL^1(s_2^1,s_1^1,\sigma^2) - GL^1(s_1^1,s_2^1,\sigma^2) + GL^1(s_2^1,s_2^1,\sigma^2)}$$

such that $\partial U(\sigma^1,\hat{\sigma}^1,\sigma^2)/\partial\sigma^1(s_1^1) \gtreqless 0$ if and only if $\hat{\sigma}^1(s_1^1) \gtreqless h(\sigma^2(s_1^2))$. Hence, if $h(\sigma^2(s_1^2)) \in (0,1)$, then $\sigma^1(s_1^1) = h(\sigma^2(s_1^2))$ is a PE. In this case, also $\sigma^1(s_1^1) = 0$ and $\sigma^1(s_1^1) = 1$ are both PEs because $\partial U(\sigma^1,\hat{\sigma}^1,\sigma^2)/\partial\sigma^1(s_1^1) > 0$ for $\hat{\sigma}^1(s_1^1) = 1$ and $\partial U(\sigma^1,\hat{\sigma}^1,\sigma^2)/\partial\sigma^1(s_1^1) < 0$ for $\hat{\sigma}^1(s_1^1) = 0$. If $h(\sigma^2(s_1^2)) > 1$, then $\partial U(\sigma^1,\hat{\sigma}^1,\sigma^2)/\partial\sigma^1(s_1^1) < 0$ for $\hat{\sigma}^1(s_1^1) \in [0,1]$. Hence, the only PE is $\sigma^1(s_1^1) = 0$. Similarly, if $h(\sigma^2(s_1^2)) < 0$, $\partial U(\sigma^1,\hat{\sigma}^1,\sigma^2)/\partial\sigma^1(s_1^1) > 0$ for $\hat{\sigma}^1(s_1^1) \in [0,1]$. Hence, the only PE is $\sigma^1(s_1^1) = 1$. Finally, by the same token, if $h(\sigma^2(s_1^2)) \in \{0,1\}$, then $\sigma^1(s_1^1) = 0$ and $\sigma^1(s_1^1) = 1$ are both PEs. The correspondence describing all PEs for player 1 is thus given by:

$$R^1(\sigma^2(s_1^2)) = \begin{cases} 0 & \text{if } h(\sigma^2(s_1^2)) > 1 \\ \{0, h(\sigma^2(s_1^2)), 1\} & \text{if } h(\sigma^2(s_1^2)) \in [0,1] \\ 1 & \text{if } h(\sigma^2(s_1^2)) < 0 \end{cases}$$

Define $\mathcal{R} = \{(\sigma^2(s_1^2), R^1(\sigma^2(s_1^2)))|\sigma^2(s_1^2) \in [0,1]\}$. In the next step, we prove that there exists a subset $\mathcal{L} \subseteq \mathcal{R}$ such that $\mathcal{L}$ is connected and includes the points $(0, R^1(0))$ and $(1, R^1(1))$. We distinguish three cases. (Case 1 is illustrated in Figure A.1.)

Case 1: Suppose $h(0) \geq 1$. Hence, $0 \in R^1(0)$. If $h(\sigma^2(s_1^2)) \geq 0$ for all $\sigma^2(s_1^2) \in [0,1]$, then $\mathcal{L} = \{(x,0)|x \in [0,1]\} \subseteq \mathcal{R}$ is connected and we are done. Otherwise, if $h(\sigma^2(s_1^2)) < 0$ for some value $\sigma^2(s_1^2) \in (0,1)$, then there exists $\sigma^2_{\text{II}} \in (0,1)$ and $\sigma^2_{\text{I}} \in [0,\sigma^2_{\text{II}})$ such that $\sigma^2_{\text{II}} = \min_{\sigma^2(s_1^2)\in(0,1)}\{\sigma^2(s_1^2)\,|h(\sigma^2(s_1^2)) = 0\}$ and $\sigma^2_{\text{I}} = \max_{\sigma^2(s_1^2)\in[0,\sigma^2_{\text{II}})}$ $\{\sigma^2(s_1^2)\,|h(\sigma^2(s_1^2)) = 1\}$. Since $h(\sigma^2(s_1^2))$ is a C^1 function, the set $\{(x,0)|x \in [0,\sigma^2_{\text{II}}]\} \cup \{(x,h(x))|x \in [\sigma^2_{\text{I}},\sigma^2_{\text{II}}]\} \cup \{(x,1)|x \in [\sigma^2_{\text{I}},\sigma^2_{\text{II}}]\} \subseteq \mathcal{R}$ is connected. If $h(\sigma^2(s_1^2)) \leq 1$ for all $\sigma^2(s_1^2) \geq \sigma^2_{\text{II}}$, then the set $\mathcal{L}$ = $\{(x,0)|x \in [0,\sigma^2_{\text{II}}]\} \cup \{(x,h(x))|x \in [\sigma^2_{\text{I}},\sigma^2_{\text{II}}]\} \cup \{(x,1)|x \in [\sigma^2_{\text{I}},1]\}$ is connected and includes the point $(0,R^1(0))$ as well as $(1,R^1(1))$—so we are done. Otherwise, if $h(\sigma^2(s_1^2)) > 1$ for some value $\sigma^2(s_1^2) \in (\sigma^2_{\text{II}},1]$, then there exists $\sigma^2_{\text{IV}} \in (\sigma^2_{\text{II}},1]$ and $\sigma^2_{\text{III}} \in (\sigma^2_{\text{II}},\sigma^2_{\text{IV}})$ such that $\sigma^2_{\text{IV}} = \min_{\sigma^2(s_1^2)\in(\sigma^2_{\text{II}},1]}\{\sigma^2(s_1^2)\,|h(\sigma^2(s_1^2)) = 1\}$ and $\sigma^2_{\text{III}} = \max_{\sigma^2(s_1^2)\in[\sigma^2_{\text{II}},\sigma^2_{\text{IV}})}\{\sigma^2(s_1^2)\,|h(\sigma^2(s_1^2)) = 0\}$. The set $\{(x,0)|x \in [0,\sigma^2_{\text{II}}]\} \cup \{(x,h(x))|x \in [\sigma^2_{\text{I}},\sigma^2_{\text{II}}]\} \cup \{(x,1)|x \in [\sigma^2_{\text{I}},\sigma^2_{\text{IV}}]\} \cup \{(x,h(x))|x \in [\sigma^2_{\text{III}},\sigma^2_{\text{IV}}]\} \cup \{(x,0)|x \in [\sigma^2_{\text{III}},\sigma^2_{\text{IV}}]\} \subseteq \mathcal{R}$ is a connected set. If $h(\sigma^2(s_1^2)) \geq 0$ for all $\sigma^2(s_1^2) \geq \sigma^2_{\text{IV}}$, the set $\mathcal{L}$ = $\{(x,0)|x \in [0,\sigma^2_{\text{II}}]\} \cup \{(x,h(x))|x \in [\sigma^2_{\text{I}},\sigma^2_{\text{II}}]\} \cup \{(x,1)|x \in [\sigma^2_{\text{I}},\sigma^2_{\text{IV}}]\} \cup \{(x,h(x))|x \in [\sigma^2_{\text{III}},\sigma^2_{\text{IV}}]\} \cup \{(x,0)|x \in [\sigma^2_{\text{III}},1]\} \subseteq \mathcal{R}$ is connected and includes the point $(0,R^1(0))$ as well as $(1,R^1(1))$—so we are done. Otherwise, if $h(\sigma^2(s_1^2)) < 0$ for some value $\sigma^2(s_1^2) \in (\sigma^2_{\text{IV}},1)$, we can proceed in the same way as we did from σ^2_{II} onward and merge sets in the same manner as before to construct a set $\mathcal{L}$ that is a connected subset of $\mathcal{R}$ including the point $(0,R^1(0))$ as well as $(1,R^1(1))$.

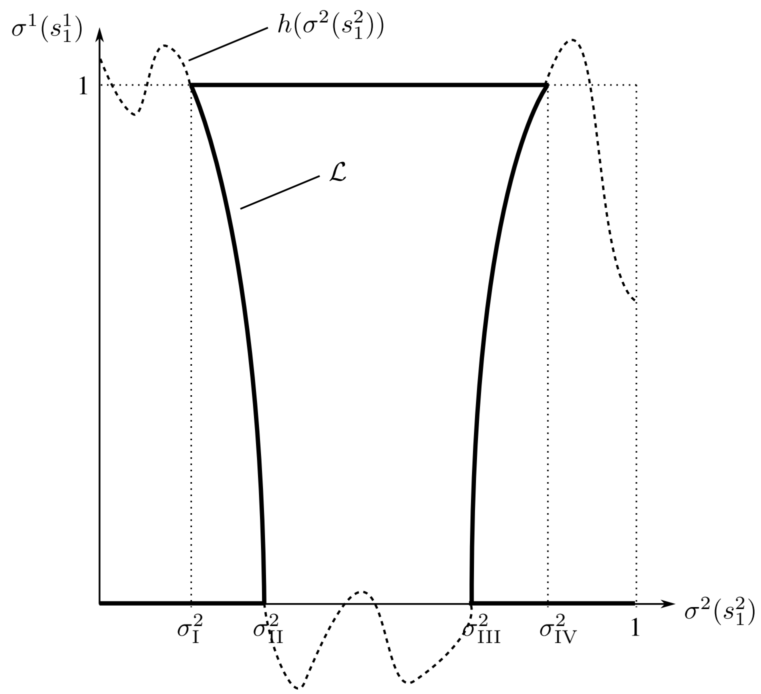

Figure A.1.: Illustration of the construction of the set $\mathcal{L}$

Case 2: Suppose $h(0) \leq 0$. The derivation of the set $\mathcal{L}$ goes along the same lines as in Case 1, starting right after σ^2_{II}.

Case 3: Suppose $h(0) \in (0, 1)$. If $h(\sigma^2(s_1^2)) \in (0, 1)$ for all $\sigma^2(s_1^2) \in [0, 1]$, then the set $\mathcal{L}$ = $\{(x, h(x))|x \in [0, 1]\} \subseteq \mathcal{R}$ is a connected set—so we are done. Otherwise, if $h(\sigma^2(s_1^2)) \geq 1$ (≤ 0) for some $\sigma^2(s_1^2) \in (0, 1]$, then the construction of the set $\mathcal{L}$ works in analogy to *Case 1* (*Case 2*).

Thus, given that s_1^1 and s_2^1 are not redundant, there always exists a connected subset $\mathcal{L} \subseteq \mathcal{R}$ including some points $(0, R^1(0))$ and $(1, R^1(1))$.

Suppose now s_1^1 and s_2^1 are redundant for some strategy $\tilde{\sigma}^2$ of player 2. For this strategy of player 2 both pure strategies of player 1 induce the same lotteries and she is indifferent between any mixture over her two pure strategies, i.e., $\mathcal{R}(\tilde{\sigma}^2(s_1^2)) = [0, 1]$. The construction of the set $\mathcal{L}$ is then analogous to the case of non-redundant strategies. For every strategy of player 2 for which the pure strategies of player 1 are redundant, however, $\mathcal{L} = [0, 1]$.

With analogous reasoning applying for player 2, the graphs $(x, R^1(x))$ and $(x, R^2(x))$ must have an intersection in $\mathbb{R}^2$. This intersection constitutes a PNE. □

Proof of Proposition 3.5. Suppose that (s^i, s^{-i}) is a Nash equilibrium in (materially) weakly dominant strategies. First, we are going to argue that a loss-averse player i has no strictly profitable deviation such that (s^i, s^{-i}) is a CPNE. Thereafter, we are going to show that any strategy profile (σ^i, σ^{-i}) in which some player does not play her (materially) weakly dominant pure strategy with probability one is not a CPNE.

As a preliminary result, we are going to establish that $U^i(s^i, s^i, \sigma^{-i}) > U^i(\tilde{s}^i, \tilde{s}^i, \sigma^{-i})$ for all $\tilde{s}^i \in \mathcal{S}^i/\{s^i\}$ and $\sigma^{-i} \in \Sigma^{-i}$. To this end, we denote by $\chi((\hat{s}^{-i}, \hat{\theta})|\sigma^{-i}) := Q(\hat{\theta})\left(\Pi_{j \neq i}\sigma^j(\hat{s}^j)\right)$ the probability that the particular combination of player i's opponents' strategy profile $\hat{s}^{-i} = (\hat{s}^j)_{j \neq i}$ and the state of the world $\hat{\theta}$ is realized. Furthermore, define $\mathcal{X} := \Sigma^{-i} \times \Theta$. Then

$$U^i(\tilde{s}^i, \tilde{s}^i, \sigma^{-i}) = \sum_{(\hat{s}^{-i}, \hat{\theta}) \in \mathcal{X}} \chi((\hat{s}^{-i}, \hat{\theta})|\sigma^{-i}) u^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta})$$
$$-\frac{\eta(\lambda-1)}{2} \sum_{(\hat{s}^{-i}, \hat{\theta}) \in \mathcal{X}} \chi((\hat{s}^{-i}, \hat{\theta})|\sigma^{-i}) \sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}} \chi((\tilde{s}^{-i}, \tilde{\theta})|\sigma^{-i}) |u^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta}) - u^i((\tilde{s}^i, \tilde{s}^{-i}), \tilde{\theta})|. \tag{A.18}$$

Defining $\mathcal{X}_+(\hat{s}^{-i}, \hat{\theta}) \equiv \{(\tilde{s}^{-i}, \tilde{\theta}) \neq (\hat{s}^{-i}, \hat{\theta}) | u^i((s^i, \hat{s}^{-i}), \hat{\theta}) \geq u^i((s^i, \tilde{s}^{-i}), \tilde{\theta})\}$ and $\mathcal{X}_-(\hat{s}^{-i}, \hat{\theta}) \equiv \{(\tilde{s}^{-i}, \tilde{\theta}) | u^i((s^i, \hat{s}^{-i}), \hat{\theta}) < u^i((s^i, \tilde{s}^{-i}), \tilde{\theta})\}$, differentiation of (A.18) yields

$$\frac{dU^i(\tilde{s}^i, \tilde{s}^i, \sigma^{-i})}{du^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta})} =$$
$$\chi((\hat{s}^{-i}, \hat{\theta})|\sigma^{-i}) \left\{ 1 - \eta(\lambda-1) \left[\sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}_+(\hat{s}^{-i}, \hat{\theta})} \chi((\tilde{s}^{-i}, \tilde{\theta})|\sigma^{-i}) - \sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}_-(\hat{s}^{-i}, \hat{\theta})} \chi((\tilde{s}^{-i}, \tilde{\theta})|\sigma^{-i}) \right] \right\}. \tag{A.19}$$

Together $\sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}_+(\hat{s}^{-i}, \hat{\theta})} \chi((\tilde{s}^{-i}, \tilde{\theta})|\sigma^{-i}) - \sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}_-(\hat{s}^{-i}, \hat{\theta})} \chi((\tilde{s}^{-i}, \tilde{\theta})|\sigma^{-i}) \leq 1 - \chi((\hat{s}^{-i}, \hat{\theta})|\sigma^{-i})$ and $\eta(\lambda - 1) \leq 1$ imply that $\frac{dU^i(\tilde{s}^i, \tilde{s}^i, \sigma^{-i})}{du^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta})} > 0$. With s^i being (materially) weakly dominant, we have $u^i((s^i, \hat{s}^{-i}), \hat{\theta}) \geq u^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta})$ for all $(\tilde{s}^i, \hat{s}^{-i}) \in \mathcal{S}^i \times \mathcal{S}^{-i}$ and $\hat{\theta} \in \Theta$, where for each $(\tilde{s}^i, \hat{s}^{-i}) \in \mathcal{S}^i/\{s^i\} \times \mathcal{S}^{-i}$ the inequality is strict for some $\hat{\theta} \in \Theta$. It then follows from (A.19) that $U^i(s^i, s^i, \sigma^{-i}) > U^i(\tilde{s}^i, \tilde{s}^i, \sigma^{-i})$.

Now, consider the Nash equilibrium in (materially) weakly dominant strategies (s^i, s^{-i}). As we showed before (by setting $\sigma^{-i} = s^{-i}$), there is no profitable pure strategy deviation for player i. Furthermore, as we established in the proof of Proposition 3.2, player i's expected utility from playing some mixed strategy σ^i is at most as large as her maximum expected utility from that mixed strategy's pure strategy components, which themselves do not constitute profitable deviations. Hence, given her opponents play their (materially) weakly dominant strategies s^{-i}, s^i is a best response for player i, such that (s^i, s^{-i}) is a CPNE.

Finally, suppose there exists some CPNE $(\tilde{\sigma}^i, \tilde{\sigma}^{-i})$ different from (s^i, s^{-i}). Since $(\tilde{\sigma}^i, \tilde{\sigma}^{-i})$ differs from (s^i, s^{-i}), there must exist some player, say player i, who does not play her (materially) weakly dominant pure strategy s^i with certainty. If player i plays some pure strategy $\tilde{s}^i \neq s^i$, then playing s^i is a strictly profitable deviation (see above). If player i plays a mixed strategy, then, for this mixture to be a CPE, she has to randomize only over pure strategies that induce the same probabilistic consequences—cf. Proposition 3.2. The probabilistic consequences of player i's (materially) weakly dominant strategy s^i, however, are unique; i.e., $L^i(s^i, \sigma^{-i}) \neq L^i(\tilde{s}^i, \sigma^{-i})$ for all $\tilde{s}^i \neq s^i$. Therefore, if player i plays a mixed strategy in the CPNE, this mixed strategy must not involve s^i. But then playing s^i is a strictly profitable deviation for player i, because, as follows from the proof of Proposition 3.2, the expected utility from playing some mixed strategy is at most as large as the maximum expected utility from that mixed strategy's pure strategy components. Thus, overall, $(\tilde{\sigma}^i, \tilde{\sigma}^{-i})$ is not a CPNE. □

Proof of Proposition 3.6. We will show that the results from Proposition 3.1, Proposition 3.2, Proposition 3.4, Corollary 3.1 and Proposition 3.5 remain to hold in turn:

Regarding Proposition 3.1 for multidimensional outcomes:

(i) The coefficients $q^i_{km}(\sigma^{-i})$ for multidimensional outcomes differ from their counterparts for one-dimensional outcomes only in the sense that every comparison of of two outcomes is replaced by a sum of possible gains and losses instead of just one gain or loss. In the same way, the coefficients $b_k(\sigma^{-i})$ only differ in the sense that the material utility from an outcome is replaced by a sum over material utilities in different dimensions. Continuity of the coefficients, however, is maintained and therefore the matrix $A(\sigma^{-i})$ for multidimensional outcomes has qualitatively identical properties to the one for one-dimensional outcomes.

Step 1 from the proof follows directly. It remains to show that non-redundancy of all pure strategies contained in $\Gamma(\bar{\sigma}^i)$ implies full rank of matrix $A'(\sigma^{-i})$—cf. Step 2—which boils down to showing that one pure strategy contained in $\Gamma(\bar{\sigma}^i)$ is redundant given that

$$\begin{aligned} & GL(\sigma^i_{\tilde{x}}, \bar{\sigma}^i, \sigma^{-i}) - GL(\bar{\sigma}^i, \bar{\sigma}^i, \sigma^{-i}) - GL(\sigma^i_{\tilde{x}}, \sigma^i_{\tilde{x}}, \sigma^{-i}) + GL(\bar{\sigma}^i, \sigma^i_{\tilde{x}}, \sigma^{-i}) = 0 \\ & \Leftrightarrow \frac{1}{2} \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|(\bar{\sigma}^i, \sigma^{-i})) \cdot P^i(\tilde{u}|(\bar{\sigma}^i, \sigma^{-i})) \cdot \sum_{r=1}^{R} |u_r - \tilde{u}_r| \\ & - \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|(\bar{\sigma}^i, \sigma^{-i})) \cdot P^i(\tilde{u}|(\sigma^i_{\tilde{x}}, \sigma^{-i})) \cdot \sum_{r=1}^{R} |u_r - \tilde{u}_r| \\ & + \frac{1}{2} \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|(\sigma^i_{\tilde{x}}, \sigma^{-i})) \cdot P^i(\tilde{u}|(\sigma^i_{\tilde{x}}, \sigma^{-i})) \cdot \sum_{r=1}^{R} |u_r - \tilde{u}_r| = 0 \end{aligned} \tag{A.20}$$

holds, which is the analogue to (A.10) for multidimensional outcomes. Let $\Lambda^i_r(u) = \{(s, \theta) \in S \times \Theta \,|\, u^i_r(s, \theta) = u\}$ denote the set of (s, θ) combinations that result in some specific payoff $u_r \in \mathcal{U}^i_r$ for player $i \in \mathcal{I}$ in dimension r. The probability of u_r being realized for player i given the strategies σ^i and σ^{-i} then is $P^i(u_r|\sigma) = \sum_{(s,\theta) \in \Lambda^i_r(u)} Q(\theta) \Pi^I_{j=1} \sigma^j(s^j)$ and (A.20) can be rewritten equivalently as

$$\sum_{r=1}^{R}\Bigg[\frac{1}{2}\sum_{u_r\in\mathcal{U}_r^i}\sum_{\tilde{u}_r\in\mathcal{U}_r^i}P^i(u_r|(\bar{\sigma}^i,\sigma^{-i}))\cdot P^i(\tilde{u}_r|(\bar{\sigma}^i,\sigma^{-i}))\cdot|u_r-\tilde{u}_r|$$
$$-\sum_{u_r\in\mathcal{U}_r^i}\sum_{\tilde{u}_r\in\mathcal{U}_r^i}P^i(u_r|(\bar{\sigma}^i,\sigma^{-i}))\cdot P^i(\tilde{u}_r|(\sigma^i_{\tilde{x}},\sigma^{-i}))\cdot|u_r-\tilde{u}_r|$$
$$+\frac{1}{2}\sum_{u_r\in\mathcal{U}_r^i}\sum_{\tilde{u}_r\in\mathcal{U}_r^i}P^i(u_r|(\sigma^i_{\tilde{x}},\sigma^{-i}))\cdot P^i(\tilde{u}_r|(\sigma^i_{\tilde{x}},\sigma^{-i}))\cdot|u_r-\tilde{u}_r|\Bigg]=0$$

By Lemma 3.1, this holds true if and only if the lotteries over material utility outcomes induced by $\bar{\sigma}^i$ and $\sigma^i_{\tilde{x}}$ are identical for every dimension $r=1,\dots,R$. Then the lottery that is induced by the pure strategy being played with zero probability under $\sigma^i_{\tilde{x}}$ and with positive probability under $\bar{\sigma}^i$ is a linear combination of the lotteries that are induced by the other pure strategies being played with positive probability for every dimension. Note that the weights of the linear combination have to be identical for every dimension since they are determined solely by $\bar{\sigma}^i$ and $\sigma^i_{\tilde{x}}$. Thus, the lottery over multidimensional outcomes induced by the pure strategy being played with probability zero under $\sigma^i_{\tilde{x}}$ is a linear combination of the lotteries that are induced by the other pure strategies being played with positive probability, implying redundancy of $\sigma^i_{\tilde{x}}$.

(ii) The matrix $A'(\sigma^{-i})$ for the case of multidimensional outcomes does not qualitatively differ from the matrix for the case of one-dimensional outcomes and hence, the proof is identical to the proof of Proposition 3.1(ii).

Regarding Proposition 3.2 for multidimensional outcomes:

The proof for multidimensional outcomes equals the proof of Proposition 3.2 up to (A.15), where multidimensionality has to be considered. Denoting the probability of u_r being realized for player i given the strategies σ^i and σ^{-i} by $P^i(u_r|\sigma^i,\sigma^{-i})$, the analogue to (A.15) for multidimensional outcomes is

$$\beta U^i(\sigma_{m'},\sigma_{m'},\sigma^{-i})+(1-\beta)U^i(\sigma_{m''},\sigma_{m''},\sigma^{-i})>$$
$$U^i(\beta\sigma_{m'}+(1-\beta)\sigma_{m''},\beta\sigma_{m'}+(1-\beta)\sigma_{m''},\sigma^{-i})$$
$$\Leftrightarrow\sum_{r=1}^{R}\Bigg[\frac{1}{2}\sum_{u_r\in\mathcal{U}_r^i}\sum_{\tilde{u}_r\in\mathcal{U}_r^i}P^i(u_r|(\sigma^i_{m''},\sigma^{-i}))\cdot P^i(\tilde{u}_r|(\sigma^i_{m''},\sigma^{-i}))\cdot|u_r-\tilde{u}_r|$$
$$-\sum_{u_r\in\mathcal{U}_r^i}\sum_{\tilde{u}_r\in\mathcal{U}_r^i}P^i(u_r|(\sigma^i_{m'},\sigma^{-i}))\cdot P^i(\tilde{u}_r|(\sigma^i_{m''},\sigma^{-i}))\cdot|u_r-\tilde{u}_r|$$
$$+\frac{1}{2}\sum_{u_r\in\mathcal{U}_r^i}\sum_{\tilde{u}_r\in\mathcal{U}_r^i}P^i(u_r|(\sigma^i_{m'},\sigma^{-i}))\cdot P^i(\tilde{u}_r|(\sigma^i_{m'},\sigma^{-i}))\cdot|u_r-\tilde{u}_r|\Bigg]<0.$$

By Lemma 3.1, this last inequality holds if and only if

$$P^i(u_r|\sigma_{m'},\sigma^{-i})\neq P^i(u_r|\sigma_{m''},\sigma^{-i})$$
$$\Leftrightarrow P^i(u_r|s_{m'},\sigma^{-i})\neq P^i(u_r|s_{m''},\sigma^{-i})$$

for some $u_r \in \mathcal{U}_r^i$. Hence, for $\bar{\sigma}^i$ to be a CPE $P^i(u_r|s_{m'}, \sigma^{-i}) = P^i(u_r|s_{m''}, \sigma^{-i})$ must hold true in every dimension $r = 1, \ldots, R$ for each outcome $u_r \in \mathcal{U}_r^i$. Overall player i is only willing to mix between two actions if they induce the same lotteries over utility vectors.

Regarding Proposition 3.4 for multidimensional outcomes:

(i) Define $U_r^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$ as the expected utility derived in dimension r from playing σ^i and having expected to play $\hat{\sigma}^i$ given σ^{-i}. $U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$ is additively separable across dimensions, i.e., $U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i}) = \sum_{r=1}^{R} U_r^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$. Hence, according to the proof of Proposition 3.4(i), $U_r^i(s^i, s^i, s^{-i}) \geq U_r^i(\sigma^i, s^i, s^{-i})$ for all $\sigma^i \in \Sigma^i$ and any $r = 1, \ldots, R$. It follows directly that $U^i(s^i, s^i, s^{-i}) \geq U^i(\sigma^i, s^i, s^{-i})$ for all $\sigma^i \in \Sigma^i$. For the reverse direction, according to the proof of Proposition 3.4(i), we have $U_r^i(s^i, \sigma^i, \sigma^{-i}) > U_r^i(\sigma^i, \sigma^i, \sigma^{-i})$ for all $\sigma^i \in \Sigma^i \setminus \{s^i\}$ and any $r = 1, \ldots, R$, which implies $U^i(s^i, \sigma^i, \sigma^{-i}) > U^i(\sigma^i, \sigma^i, \sigma^{-i})$ for all $\sigma^i \in \Sigma^i \setminus \{s^i\}$.

(ii) The proof is identical to the proof of Proposition 3.4(ii).

Regarding Corollary 3.1 for multidimensional outcomes:

(i) The result follows directly from the fact that Proposition 3.2 continues to hold for multi-dimensional payoffs.

(ii) The result follows from Corollary 3.1(i) together with Proposition 3.7(i).

Regarding Proposition 3.5 for multidimensional outcomes:

According to the proof of Proposition 3.5, $U_r^i(s^i, s^i, s^{-i}) \geq U_r^i(\sigma^i, \sigma^i, s^{-i})$ for all $\sigma^i \in \Sigma^i$ and any $r = 1, \ldots, R$. It follows directly that $U^i(s^i, s^i, s^{-i}) \geq U^i(\sigma^i, \sigma^i, s^{-i})$ for all $\sigma^i \in \Sigma^i$. For the reverse direction, according to the proof of Proposition 3.5 we have $U_r^i(s^i, s^i, \sigma^{-i}) > U_r^i(\sigma^i, \sigma^i, \sigma^{-i})$ for all $\sigma^i \in \Sigma^i \setminus \{s^i\}$ and any $r = 1, \ldots, R$, which implies $U^i(s^i, s^i, \sigma^{-i}) > U^i(\sigma^i, \sigma^i, \sigma^{-i})$ for all $\sigma^i \in \Sigma^i \setminus \{s^i\}$.

$\square$

Proof of Proposition 3.7. (i) The proof is identical to the corresponding proof of Proposition 3.3.

(ii) Suppose pure strategy s_k^i is a Nash best response to s^{-i}. A deviation to any strategy profile $\sigma^i \in \Sigma^i$ yields a weakly lower expected material utility. In addition, it creates possible gains and losses, where the overall size of losses dominates the overall size of gains. With losses looming larger than gains, no deviation from a Nash best response can be profitable for a loss-averse player.

(iii) Suppose that for each $\tilde{s}^i \in \mathcal{S}^i \setminus \{s^i\}$, where $i \in \mathcal{I}$, there exists $r^i(\tilde{s}^i)$ such that $u_{r^i(\tilde{s}^i)}^i((\tilde{s}^i, s^{-i}), \tilde{\theta}) < u_{r^i(\tilde{s}^i)}^i((s^i, s^{-i}), \tilde{\theta})$. For λ sufficiently large, the impact of the loss in dimension $r^i(\tilde{s}^i)$ caused by the unilateral deviation from s^i to $\tilde{s}^i$ dominates possible gains in other dimensions and a potentially higher material utility, such that $U^i(s^i, s^i, s^{-i}) \geq U^i(\tilde{s}^i, s^i, s^{-i})$ for all $\tilde{s}^i \in \mathcal{S}^i$ holds for all players $i \in \mathcal{I}$. Therefore s can be implemented in a PNE for λ sufficiently large.

$\square$

A.4. Appendix to Chapter IV

Proof of Lemma 1. First, we establish the comparative static results listed in the lemma. Given that $e_i^*(e_j)$ is a local maximizer of $U^i(e_i, e_i, e_j)$, we have $\partial^2 U^i(e_i^*(e_j), e_i^*(e_j), e_j)/\partial e_i^2 < 0$. Implicit differentiation of the condition $\partial U^i(e_i^*(e_j), e_i^*(e_j), e_j)/\partial e_i = 0$ then yields

$$\frac{de_i^*(e_j)}{de_j} = -\frac{\partial^2 U^i(e_i^*(e_j), e_i^*(e_j), e_j)/\partial e_i \partial e_j}{\partial^2 U^i(e_i^*(e_j), e_i^*(e_j), e_j)/\partial e_i^2} = \frac{\frac{\Delta}{2}\eta(\lambda-1)}{\partial^2 U^i(e_i^*(e_j), e_i^*(e_j), e_j)/\partial e_i^2} < 0 \quad \text{(A.21)}$$

and

$$\frac{d^2 e_i^*(e_j)}{de_j^2} = \frac{\frac{\Delta}{2}\eta(\lambda-1)c'''(e_i^*(e_j))\frac{de_i^*(e_j)}{de_j}}{[\partial^2 U^i(e_i^*(e_j), e_i^*(e_j), e_j)/\partial e_i^2]^2} < 0. \quad \text{(A.22)}$$

Next, we derive player i's best response function $BR_i(e_j)$ for each of the following cases: (i) $\eta(\lambda-1) < 1$; (ii) $1 \le \eta(\lambda-1) \le \frac{2}{\Delta}c''(0)$; (iii) $1 \le \frac{2}{\Delta}c''(0) < \eta(\lambda-1)$; (iv) $\frac{2}{\Delta}c''(0) < 1 \le \eta(\lambda-1)$

Case (i): $\boldsymbol{\eta(\lambda-1) < 1}$

Recall that the derivatives of the expected utility function are given by

$$\frac{\partial U^i(e_i, e_i, e_j)}{\partial e_i} = \frac{\Delta}{2}\left[1 + \eta(\lambda-1)(e_i - e_j)\right] - c'(e_i) \quad \text{(A.23)}$$

and

$$\frac{\partial^2 U^i(e_i, e_i, e_j)}{\partial e_i^2} = \frac{\Delta}{2}\eta(\lambda-1) - c''(e_i). \quad \text{(A.24)}$$

As $\eta(\lambda-1) < 1$, we have that $\partial U^i(0,0,e_j)/\partial e_i > 0$, i.e., $U^i(e_i, e_i, e_j)$ is strictly increasing for small values of e_i irrespective of e_j. As explained in the text, this implies that $U^i(e_i, e_i, e_j)$ has an interior local maximum that is also its global maximizer. Hence, $BR_i(e_j) = e_i^*(e_j)$ for all $e_j \in [0,1]$. ||

Case (ii): $\boldsymbol{1 \le \eta(\lambda-1) \le \frac{2}{\Delta}c''(0)}$

By $c''' > 0$ and $\partial^2 U^i(0,0,e_j)/\partial e_i^2 \le 0$, we have $\partial^2 U^i(e_i, e_i, e_j)/\partial e_i^2 < 0$ for all $e_i \in (0,1]$; i.e., expected utility is strictly concave. First, consider $e_j < \frac{1}{\eta(\lambda-1)}$ such that $\partial U^i(0,0,e_j)/\partial e_i > 0$. In this case, there exists a unique value $e_i^*(e_j) \in (0,1)$ such that $\partial U^i(e_i^*(e_j), e_i^*(e_j), e_j) = 0$. By strict concavity of $U^i(e_i, e_i, e_j)$, $e_i^*(e_j)$ is the global maximizer of $U^i(e_i, e_i, e_j)$. For all $e_j \ge \frac{1}{\eta(\lambda-1)}$, we have $\partial U^i(0,0,e_j)/\partial e_i \le 0$. By strict concavity of $U^i(e_i, e_i, e_j)$, we then have $\partial U^i(e_i, e_i, e_j)/\partial e_i < 0$ for all $e_i \in (0,1]$ and the global maximizer of $U^i(e_i, e_i, e_j)$ is given by $e_i = 0$. Thus, $\bar{e} = \frac{1}{\eta(\lambda-1)}$. ||

For cases (iii) and (iv), define the following functions to make the proof more tractable,

$$\theta(e_i) = c'(e_i) - \frac{\Delta}{2}\eta(\lambda-1)e_i \qquad \text{and} \qquad \psi(e_j) = \frac{\Delta}{2}[1 - \eta(\lambda-1)e_j], \quad \text{(A.25)}$$

such that

$$\frac{\partial U^i(e_i, e_i, e_j)}{\partial e_i} = \psi(e_j) - \theta(e_i) \qquad \text{and} \qquad \frac{\partial^2 U^i(e_i, e_i, e_j)}{\partial e_i^2} = -\theta'(e_i). \tag{A.26}$$

Regarding the function $\psi(e_j)$, note that $\psi(0) > 0$, $\psi'(e_j) < 0$, and $\frac{d\psi(1)}{d\eta(\lambda-1)} = -\frac{\Delta}{2} < 0$. Furthermore, $\psi(1) \lesseqgtr 0$ if and only if $\eta(\lambda - 1) \gtreqless 1$.

Concerning the function $\theta(e_i)$, first note that $\theta(0) = 0$. Furthermore, $\theta''(e_i) > 0$, i.e., $\theta(e_i)$ is strictly convex, and $\theta'(0) \lesseqgtr 0$ if and only if $\eta(\lambda - 1) \gtreqless \frac{2}{\Delta}c''(0)$. Finally, with $\lim_{e_i\to 1} c'(e_i) = \infty$ and $\lim_{e_i\to 1} c''(e_i) = \infty$, we have $\lim_{e_i\to 1} \theta(e_i) = \infty$ and $\lim_{e_i\to 1} \theta'(e_i) = \infty$. Hence, if $\eta(\lambda - 1) > \frac{2}{\Delta}c''(0)$, the global minimizer $e_{min} \in (0, 1)$ of $\theta(e_i)$ is implicitly defined by $\theta'(e_{min}) = c''(e_{min}) - \frac{\Delta}{2}\eta(\lambda - 1) = 0$, in which case $\theta'(e_i) \lesseqgtr 0$ if and only if $e_i \lesseqgtr e_{min}$. Also, $\frac{de_{min}}{d\eta(\lambda-1)} = \frac{\Delta}{2c'''(e_{min})} > 0$, $\lim_{\eta(\lambda-1)\to\infty} e_{min} = 1$, and $\frac{d\theta(e_{min})}{d\eta(\lambda-1)} = -\frac{\Delta}{2}e_{min} < 0$. Note that e_{min} is an inflection point such that $U^i(e_i, e_i, e_j)$ is strictly convex for $e_i < e_{min}$ and strictly concave for $e_i > e_{min}$. See Figure A.2 for a graphical representation of $\theta(e_i)$ and $\psi(e_j)$.

Case (iii): $1 \leq \frac{2}{\Delta}c''(0) < \eta(\lambda - 1)$

In this case, $\psi(1) < 0$. With $\lim_{\eta(\lambda-1)\to\frac{2}{\Delta}c''(0)} \theta(e_{min}) = \theta(0) = 0$ and $\frac{d\psi(1)}{d\eta(\lambda-1)} \leq \frac{d\theta(e_{min})}{d\eta(\lambda-1)} < 0$, we have $\psi(1) < \theta(e_{min}) < 0 < \psi(0)$. Hence, there exist $\underline{e}_j$ and $\bar{e}_j$, where $[\underline{e}_j, \bar{e}_j] \subset (0, 1)$, implicitly defined by $\psi(\underline{e}_j) = 0$ and $\psi(\bar{e}_j) = \theta(e_{min})$. For each $e_j \in [\underline{e}_j, \bar{e}_j)$ there exist two values of e_i, one strictly smaller and the other strictly larger than e_{min}, such that $\psi(e_j) = \theta(e_i)$ or, equivalently, $\partial U^i(e_i, e_i, e_j)/\partial e_i = 0$. The smaller of these e_i values is a local minimizer and the larger one, denoted by $e_i^*(e_j)$, is a local maximizer of $U^i(e_i, e_i, e_j)$. Expected utility from exerting effort $e_i^*(e_j)$ amounts to

$$U^i(e_i^*(e_j), e_i^*(e_j), e_j) = U^i(0, 0, e_j) + \int_0^{e_i^*(e_j)} [\psi(e_j) - \theta(e_i)]de_i. \tag{A.27}$$

With $U^i(e_i^*(\underline{e}_j), e_i^*(\underline{e}_j), \underline{e}_j) > U^i(0, 0, \underline{e}_j)$, $U^i(e_i^*(\bar{e}_j), e_i^*(\bar{e}_j), \bar{e}_j) < U^i(0, 0, \bar{e}_j)$, and

$$\frac{d[U^i(0, 0, e_j) - U^i(e_i^*(e_j), e_i^*(e_j), e_j)]}{de_j} = -\int_0^{e_i^*(e_j)} \psi'(e_j)de_i = \frac{\Delta}{2}\eta(\lambda - 1) > 0, \tag{A.28}$$

there exists $\bar{e} \in (\underline{e}_j, \bar{e}_j)$ such that the global maximizer of $U^i(e_i, e_i, e_j)$ is $e_i^*(e_j) > 0$ for $e_j \leq \bar{e}$ and $e_i = 0$ for $e_j \geq \bar{e}$. ||

Case (iv): $\frac{2}{\Delta}c''(0) < 1 \leq \eta(\lambda - 1)$

Again, $U^i(e_i, e_i, e_j)$ is convex for $e_i < e_{min}$ and concave for $e_i > e_{min}$. First, consider $e_j \in [0, \frac{1}{\eta(\lambda-1)})$. In this case, with $\psi(e_j) > 0$, $U^i(e_i, e_i, e_j)$ is strictly increasing for small values of e_i; i.e., $\partial U^i(0, 0, e_j)/\partial e_i > 0$. There exists a unique value of e_i, denoted by $e_i^*(e_j)$ and strictly smaller than 1, such that $\psi(e_j) = \theta(e_i^*(e_j))$. With $\psi(e_j) \gtrless \theta(e_i)$—or, equivalently, $\partial U^i(e_i, e_i, e_j)/\partial e_i \gtrless 0$—if and only if $e_i \lessgtr e_i^*(e_j)$, $e_i^*(e_j)$ is the global maximizer of $U^i(e_i, e_i, e_j)$.

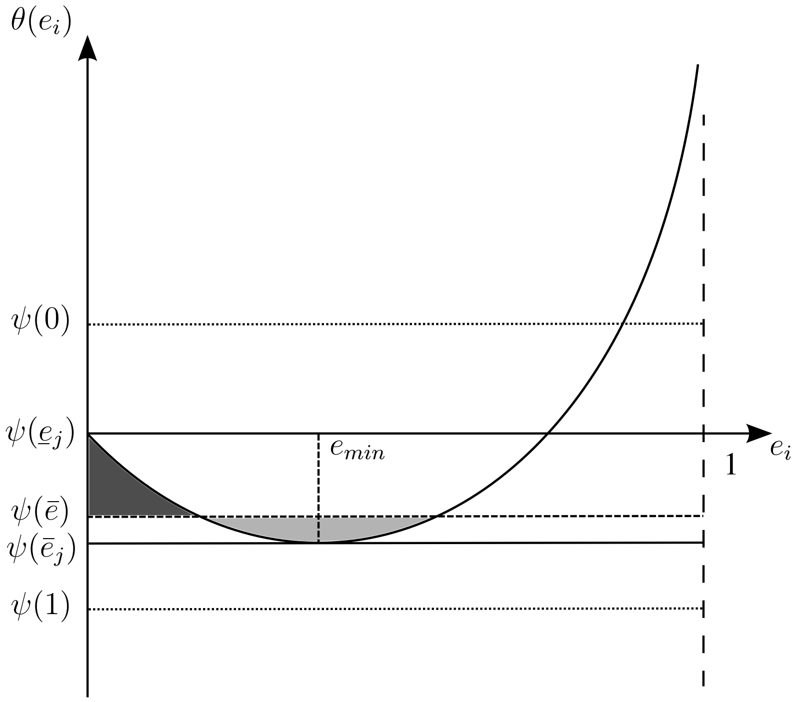

Figure A.2.: Graphical representation of the derivation of $\bar{e}$

Next, consider $e_j \in [\frac{1}{\eta(\lambda-1)}, 1]$. Note that for $e_j > \frac{1}{\eta(\lambda-1)}$, $U^i(e_i, e_i, e_j)$ is strictly decreasing for small values of e_i—i.e., $\partial U^i(0,0,e_j)/\partial e_i < 0$. For $\eta(\lambda-1) = 1$, $\psi(1) = 0$ such that $\psi(1) - \theta(e_{min}) > 0$. For $\eta(\lambda-1) \to \infty$, $e_{min} \to 1$ such that $\lim_{\eta(\lambda-1)\to\infty} \psi(1) - \theta(e_{min}) = -\infty$. From

$$\frac{d[\psi(1) - \theta(e_{min})]}{d\eta(\lambda-1)} = -\frac{\Delta}{2}(1 - e_{min}) < 0 \tag{A.29}$$

then follows the existence of a threshold $\chi' > 1$, such that $\psi(1) \lesseqgtr \theta(e_{min})$ if and only if $\eta(\lambda-1) \gtreqless \chi'$. For $\eta(\lambda-1) < \chi'$, there are two values of e_i, one strictly smaller and the other strictly larger than e_{min}, such that $\psi(1) = \theta(e_i)$. The larger of these e_i values, which we denote by $e_i^*(1)$, is strictly smaller than 1 and a local maximizer of $U^i(e_i, e_i, 1)$. With

$$U^i(e_i^*(1), e_i^*(1), 1) = U^i(0,0,1) + \int_0^{e_i^*(1)} [\psi(1) - \theta(e_i)]de_i, \tag{A.30}$$

and $U^i(0,0,1) = u(L)$, we obtain

$$\frac{dU^i(0,0,1) - U^i(e_i^*(1), e_i^*(1), 1)}{d\eta(\lambda-1)} = \frac{\Delta}{2}\int_0^{e_i^*(1)} (1 - e_i)de_i > 0, \tag{A.31}$$

where we made use of $\psi(1) - \theta(e_i^*(1)) = 0$. Furthermore, for $\eta(\lambda-1) = 1$ we have $\psi(1) = 0$ such that $U^i(0,0,1) - U^i(e_i^*(1), e_i^*(1), 1) = -\int_0^{e_i^*(1)}[-\theta(e_i)]de_i < 0$. For $\eta(\lambda-1) = \chi'$, on the other hand, we have $U^i(0,0,1) - U^i(e_i^*(1), e_i^*(1), 1) = -\int_0^{e_i^*(1)}[\theta(e_{min}) - \theta(e_i)]de_i > 0$. Hence, by the intermediate value theorem, there exists $\chi'' \in (1, \chi')$ such that $U^i(e_i^*(1), e_i^*(1), 1) \lesseqgtr U^i(0,0,1)$ if and only if $\eta(\lambda-1) \gtreqless \chi''$.

By (A.28), we have $\frac{d[U^i(0,0,e_j)-U^i(e_i^*(e_j),e_i^*(e_j),e_j)]}{de_j} > 0$, where $e_i^*(e_j)$ is defined as before. Hence, if $\eta(\lambda-1) < \chi''$ and, thus, $U^i(e_i^*(1), e_i^*(1), 1) > U^i(0,0,1)$, then $U^i(e_i^*(e_j), e_i^*(e_j), e_j) > U^i(0,0,e_j)$ for all $e_j \in [0,1]$. In this case, $BR_i(e_j) = e_i^*(e_j)$ for all $e_j \in [0,1]$. If, on the other hand, $\eta(\lambda-1) \geq \chi''$, there exists $\bar{e} \in (0,1]$ such that $e_i^*(e_j)$ is the best response to $e_j \in [0,\bar{e}]$ and $e_i = 0$ is the best response to $e_j \in [\bar{e}, 1]$. || □

Proof of Proposition 4.1. A symmetric equilibrium must be interior. First, $(e_A, e_B) = (0,0)$ cannot constitute a CPNE since $\frac{\partial U^i(e_i,e_i,0)}{\partial e_i}|_{e_i=0} > 0$. Analogously $(e_A, e_B) = (1,1)$ is not a CPNE since $\frac{\partial U^i(e_i,e_i,1)}{\partial e_i}|_{e_i=1} < 0$. Hence, a symmetric CPNE must be characterized by a solution to the first-order condition, which (given $e_i = e_j$) boils down to

$$c'(e^*) = \frac{\Delta}{2} \Leftrightarrow e^* = c'^{-1}\left(\frac{\Delta}{2}\right) = e^{NE}. \tag{A.32}$$

(i) The symmetric CPNE exists if and only if the best response to $e_j = e^{NE}$ is given by $e_i^*(e_j)$. Then the best response curves of player i and j in the symmetric equilibrium are both decreasing with identical slope. The symmetric equilibrium therefore is asymptotically unstable if and only if

$$\left.\frac{de_i^*(e_j)}{de_j}\right|_{e_j=e^{NE}} < -1 \Leftrightarrow \eta(\lambda-1) > \frac{c''(e^{NE})}{\Delta}. \tag{A.33}$$

(ii) A symmetric equilibrium cannot exist if $U^i(e_i, e_i, e^{NE})$ is strictly convex at $e_i = e^{NE}$, i.e.

$$\left.\frac{\partial^2 U^i(e_i,e_i,e_j)}{\partial e_i^2}\right|_{e_i=e^{NE}} \gtreqless 0 \Leftrightarrow \eta(\lambda-1) \gtreqless \frac{2c''(e^{NE})}{\Delta}. \tag{A.34}$$

□

Proof of Proposition 4.2. According to the proof of Proposition 4.1 (ii), a finite value of $\eta(\lambda-1)$ exists for which the symmetric CPNE ceases to exist. For any value of $\eta(\lambda-1)$ above this threshold, the two asymmetric CPNE $(e_A^{CPNE}, e_B^{CPNE}) = (e^*(0), 0)$ and $(e_A^{CPNE}, e_B^{CPNE}) = (0, e^*(0))$ exist. To see this, note that for a symmetric CPNE not to exist, each agent's best response function must display a downward discontinuity $\bar{e} \in (0,1)$ with $e_i^*(\bar{e}) = e_j^*(\bar{e}) > \bar{e}$. Since $e_i^*(e_j)$ is decreasing in e_j we conclude that $BR_i(0) = e_i^*(0) > e_i^*(\bar{e}) > \bar{e}$, such that $BR_j(e_i^*(0)) = 0$. Overall, given agent i plays his best response to zero effort, exerting zero effort is indeed a best response for agent j. The considered asymmetric CPNEs are asymptotically stable because $BR_i(e_j) = 0$ for $e_j \in [\bar{e}, 1]$ and $\bar{e} < BR_j(0) < 1$. □

Proof of Proposition 4.3. As outlined in the text, for a given effort level e_j of the opponent, agent i's best response is either minimum effort $e_i = 0$ or (in case it exists) the interior local maximizer $e_i = e_i^*(e_j)$. Furthermore, given $e_j = e^{NE}$, $e_i = e^{NE}$ always satisfies $\partial U^i(e^{NE}, e^{NE}, e^{NE})/\partial e_i = 0$. Finally, as established in the proof of Proposition 4.1,

$$\frac{\partial^2 U^i(e^{NE}, e^{NE}, e^{NE})}{\partial e_i^2} \gtreqless 0 \Leftrightarrow \eta(\lambda - 1) \gtreqless \frac{2c''(e^{NE})}{\Delta}. \tag{A.35}$$

Finally, we have $\frac{2}{\Delta}c''(e^{NE}) > 1$, which follows from $c'(0) = 0$, $c'' > 0$, and $c''' > 0$ together with $c'(e^{NE}) = \frac{\Delta}{2}$ and $e^{NE} < 1$.

For $\eta(\lambda - 1) \leq 1$, we know that $e_i = e^{NE}$ constitutes not only a local but also the global maximum of $U^i(e_i, e_i, e^{NE})$: Using the notation from Lemma 1, with $\psi(e^{NE}) > 0$, we have $\theta(e_i) \gtreqless \psi(e^{NE})$ if and only if $e_i \gtreqless e^{NE}$. In consequence, $\partial U^i(e_i, e_i, e^{NE})/\partial e_i \gtreqless 0$ if and only if $e_i \lesseqgtr e^{NE}$, such that $e_i = e^{NE}$ constitutes the global optimizer of $U^i(e_i, e_i, e^{NE})$. In particular, note that $U^i(e^{NE}, e^{NE}, e^{NE}) > U^i(0, 0, e^{NE})$.

For $\eta(\lambda - 1) \geq \frac{2c''(e^{NE})}{\Delta}$, on the other hand, $e_i = e^{NE}$ is not a candidate for agent i's best response to $e_j = e^{NE}$: With $\partial^2 U^i(e^{NE}, e^{NE}, e^{NE})/\partial e_i^2 \geq 0$, $e_i = e^{NE}$ is either a local minimum or an interior inflection point of a strictly decreasing function. In either case, since $\partial^2 U^i(0, 0, e^{NE})/\partial e_i^2 > 0$, it follows that $\partial U^i(0, 0, e^{NE})/\partial e_i < 0$, such that $U^i(e^{NE}, e^{NE}, e^{NE}) < U^i(0, 0, e^{NE})$.

To conclude the argument, note that

$$\frac{\partial \left[U^i(e^{NE}, e^{NE}, e^{NE}) - U^i(0, 0, e^{NE})\right]}{\partial \eta(\lambda - 1)} = -\frac{\Delta e^{NE^2}}{4} < 0, \tag{A.36}$$

i.e., as $\eta(\lambda - 1)$ increases, $U^i(e^{NE}, e^{NE}, e^{NE}) - U^i(0, 0, e^{NE})$ monotonically decreases at a constant rate. By the intermediate value theorem, there exists $\overline{\chi} \in (1, 2c''(e^{NE})/\Delta)$ such that $U^i(e^{NE}, e^{NE}, e^{NE}) \geq U^i(0, 0, e^{NE})$ if and only if $\eta(\lambda - 1) \leq \overline{\chi}$. For $\eta(\lambda - 1) < 2c''(e^{NE})/\Delta$, we have $\partial^2 U^i(e^{NE}, e^{NE}, e^{NE})/\partial e_i^2 < 0$. Hence, for $\eta(\lambda - 1) \in (1, 2c''(e^{NE})/\Delta)$, $e_i = e^{NE}$ is a local maximizer of $U^i(e_i, e_i, e^{NE})$ and thus, next to $e_i = 0$, the only candidate for agent i's best response to $e_j = e^{NE}$. Thus, the symmetric CPNE exists if and only if $\eta(\lambda - 1) \leq \overline{\chi}$. $\square$

Proof of Lemma 2. The fact that the set $\Theta_i^{PE}(e_j)$ is an interval is established in the text. It thus remains to establish the comparative statics with regard to the boundaries of this set. Differentiation of $\underline{\theta}(\bar{e}_i(e_j)) = \psi(e_j)$ with respect to e_j yields

$$\frac{d\bar{e}_i(e_j)}{de_j}\underline{\theta}'(\bar{e}_i(e_j)) = \psi'(e_j) < 0 \tag{A.37}$$

and

$$\frac{d^2\bar{e}_i(e_j)}{de_j^2}\underline{\theta}'(\bar{e}_i(e_j)) + \left(\frac{d\bar{e}_i(e_j)}{de_j}\right)^2 \underline{\theta}''(\bar{e}_i(e_j)) = \psi''(e_j) = 0 \tag{A.38}$$

Due to the convexity of $\underline{\theta}$, $\underline{\theta}(0) = 0$, and $\psi(e_j) > 0$, we must have $\underline{\theta}'(\bar{e}_i(e_j)) > 0$, such that (A.37) implies $\frac{d\bar{e}_i(e_j)}{de_j} < 0$. Since furthermore $\underline{\theta}''(\bar{e}_i(e_j)) > 0$, (A.38) implies $\frac{d^2\bar{e}_i(e_j)}{de_j^2} < 0$.

By analogous reasoning, differentiation of $\bar{\theta}(\underline{e}_i(e_j)) = \psi(e_j)$ with respect to e_j yields $\frac{d\underline{e}_i(e_j)}{de_j} < 0$ and $\frac{d^2\underline{e}_i(e_j)}{de_j^2} < 0$ □

Proof of Proposition 4.5. According to Proposition 4.3, a symmetric CPNE exists if and only if $\eta(\lambda - 1) \leq \bar{\chi}$. This symmetric CPNE is identical to the Nash equilibrium and, thus, also a PNE and a PPNE. It remains to show that there exists $\tilde{\chi} > \bar{\chi}$ such that the same holds true for $\eta(\lambda - 1) \in (\bar{\chi}, \tilde{\chi}]$. For this purpose, note that the smallest effort level that is a PE is strictly positive. For $\eta(\lambda-1) = \bar{\chi}$, we have that $U^i(e^{NE}, e^{NE}, e^{NE}) = U^i(0, 0, e^{NE}) > U^i(e_i, e_i, e^{NE})$ for all $e \notin \{0, e^{NE}\}$. As playing $e_i = 0$ is not a PE, (e^{NE}, e^{NE}) is the unique PPNE in this case. For $\eta(\lambda - 1)$ marginally larger than $\bar{\chi}$, $e_i = e^{NE}$ remains a local maximum of $U^i(e_i, e_i, e^{NE})$ and only effort levels marginally close to zero provide a higher utility when the player expects to play them than expecting to play and playing e^{NE}. As the smallest effort that constitutes a PE is strictly larger than 0, however, there is no PE that provides higher utility than expecting and playing $e_i = e^{NE}$. As a consequence, there exists $\tilde{\chi} > \bar{\chi}$ such that $(e_i, e_j) = (e^{NE}, e^{NE})$ is a PPNE for all $\eta(\lambda - 1) \in [\bar{\chi}, \tilde{\chi}]$. □

Bibliography

ABBINK, K., AND B. HERRMANN (2011): "The moral costs of nastiness," *Economic Inquiry*, 49(2), 631–633.

ABELER, J., A. FALK, L. GOETTE, AND D. HUFFMAN (2011): "Reference Points and Effort Provision," *American Economic Review*, 101, 470–492.

ADAMS, R. B., AND P. FUNK (2012): "Beyond the Glass Ceiling: Does Gender Matter?," *Management Science*, 58(2), 219–235.

AMEGASHIE, J. A. (2013): "Sabotage in contests: an overview," working paper.

ANDREONI, J., AND L. VESTERLUND (2001): "Which is the Fair Sex? Gender Differences in Altruism," *The Quarterly Journal of Economics*, 116(1), 293–312.

ANTONOVICS, K., P. ARCIDIACONO, AND R. WALSH (2009): "The Effects of Gender Interactions in the Lab and in the Field," *The Review of Economics and Statistics*, 91(1), 152–162.

APESTEGUIA, J., G. AZMAT, AND N. IRIBERRI (2012): "The Impact of Gender Composition on Team Performance and Decision Making: Evidence from the Field," *Management Science*, 58(1), 78–93.

AUMANN, R. J. (1985): "What is Game Theory Trying to Accomplish," *Frontiers of Economics*.

AUMANN, R. J., AND A. BRANDENBURGER (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63, 1161–1180.

BALAFOUTAS, L., F. LINDNER, AND M. SUTTER (2012): "Sabotage in tournaments: Evidence from a natural experiment," *Kyklos*, 65(4), 425–441.

BARANKAY, I. (2012): "Rank incentives: Evidence from a randomized workplace experiment," working paper.

BARBER, B. M., AND T. ODEAN (2001): "Boys will be boys: Gender, overconfidence, and common stock investment," *The Quarterly Journal of Economics*, 116(1), 261–292.

BELICA, M., AND K.-M. EHRHART (2014): "Personal Equilibrium in Bayesian Games," working paper.

BELL, D. (1985): "Disappointment in Decision Making under Uncertainty," *Operations Research*, 33, 1–27.

BENGTSSON, C., M. PERSSON, AND P. WILLENHAG (2005): "Gender and overconfidence," *Economics Letters*, 86(2), 199–203.

BERGERHOFF, J., AND A. VOSEN (2014): "How Being Behind Can Get You Ahead," mimeo.

BEYER, S., AND E. M. BOWDEN (1997): "Gender Differences in Self-Perceptions: Convergent Evidence from Three Measures of Accuracy and Bias," *Personality and Social Psychology Bulletin*, 23(2), 157–172.

BLANCO, M., D. ENGELMANN, A. K. KOCH, AND H. NORMANN (2010): "Belief elicitation in experiments: is there a hedging problem?," *Experimental Economics*, 13(4), 412–438.

BRANDENBURGER, A., AND B. NALEBUFF (1998): *Co-opetition*. Crown Business.

BUSER, T. (2012): "The Impact of the Menstrual Cycle and Hormonal Contraceptives on Competitiveness," *Journal of Economic Behavior & Organization*, 83(1), 1–10.

CARPENTER, J., P. H. MATTHEWS, AND J. SCHIRM (2010): "Tournaments and Office Politics: Evidence from a Real Effort Experiment," *American Economic Review*, 100(1), 504–17.

CHARNESS, G., AND U. GNEEZY (2012): "Strong evidence for gender differences in risk taking," *Journal of Economic Behavior & Organization*, 83(1), 50–58.

CHARNESS, G., D. MASCLET, AND M. C. VILLEVAL (2013): "The dark side of competition for status," *Management Science*, 60(1), 38–55.

CHEN, K. P. (2003): "Sabotage in Promotion Tournaments," *Journal of Law, Economics, and Organization*, 19(1), 119–140.

CHOWDHURY, M., AND O. GÜRTLER (2013): "Sabotage in Contests: A Survey," working paper.

COTTON, C., F. MCINTYRE, AND J. PRICE (2011): "Gender Differences in Competition: A Theoretical Assessment of the Evidence," working paper.

——— (2013): "Gender Differences in Repeated Competition: Evidence from School Math Contests," *Journal of Economic Behavior & Organization*, 86(0), 52–66.

CRAWFORD, V. P., AND J. MENG (2011): "New York City Cab Drivers' Labor Supply Revisited: Reference-Dependent Preferences with Rational-Expectations Targets for Hours and Income," *American Economic Review*, 101, 1912–1932.

CROCKER, J., R. K. LUHTANEN, M. L. COOPER, AND A. BOUVRETTE (2003): "Contingencies of Self-Worth in College Students: Theory and Measurement.," *Journal of Personality and Social Psychology*, 85(5), 894.

CROSON, R., AND U. GNEEZY (2009): "Gender Differences in Preferences," *Journal of Economic Literature*, 47(2), 448–474.

CRUICKSHANK, D. (2000): "Competition in UK Banking: A Report to the Chancellor of the Exchequer," London: HM Treasury: `http://www.hm-treasury.gov.uk/d/BankReviewChaptfour.pdf`.

DAHREMÖLLER, C. (2013): "Unshrouding for competitive advantage," *Journal of Economics & Management Strategy*, 22(3), 551–568.

DAIDO, K., AND H. ITOH (2007): "The Pygmalion and Galatea Effects: An Agency Model with Reference-Dependent Preferences and Applications to Self-Fulfilling Prophecy," working paper.

DAIDO, K., K. MORITA, T. MUROOKA, AND H. OGAWA (2013): "Task Assignment Under Agent Loss Aversion," *Economics Letters*, 121, 35–38.

DAIDO, K., AND T. MUROOKA (2014): "Team Incentives and Reference-Dependent Preferences," working paper.

DARGNIES, M. P. (2012): "Men Too Sometimes Shy Away from Competition: The Case of Team Competition," *Management Science*, 58(11), 1982–2000.

DATO, S., A. GRUNEWALD, AND D. MÜLLER (2015a): "Expectation-Based Loss Aversion and Rank-Order Tournaments," working paper.

DATO, S., A. GRUNEWALD, AND D. MÜLLER (2015b): "Expectation-Based Loss Aversion and Strategic Interaction," mimeo.

DATTA GUPTA, N., A. POULSEN, AND M. C. VILLEVAL (2013): "Gender Matching and Competitiveness: Experimental Evidence," *Economic Inquiry*, 51(1), 816–835.

DECHENAUX, E., D. KOVENOCK, AND R. SHEREMETA (2012): "A Survey of Experimental Research on Contests, All-Pay Auctions and Tournaments," working paper.

DEL CORRAL, J., J. PRIETO-RODRIGUEZ, AND R. SIMMONS (2010): "The effect of incentives on sabotage: the case of Spanish football," *Journal of Sports Economics*, 11(3), 243–260.

DELLAVIGNA, S. (2009): "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47(2), 315–372.

DEMOUGIN, D., AND C. FLUET (2003): "Inequity Aversion in Tournaments," working paper.

DEUTSCHER, C., B. FRICK, O. GÜRTLER, AND J. PRINZ (2013): "Sabotage in Tournaments with Heterogeneous Contestants: Empirical Evidence from the Soccer Pitch," *The Scandinavian Journal of Economics*, 115(4), 1138–1157.

DOHMEN, T., AND A. FALK (2011): "Performance Pay and Multidimensional Sorting: Productivity, Preferences, and Gender," *American Economic Review*, 101(2), 556–90.

DOHMEN, T., A. FALK, K. FLIESSBACH, U. SUNDE, AND B. WEBER (2011): "Relative versus Absolute Income, Joy of Winning, and Gender: Brain Imaging Evidence," *Journal of Public Economics*, 95(3), 279 – 285.

DOHMEN, T., A. FALK, D. HUFFMAN, U. SUNDE, J. SCHUPP, AND G. G. WAGNER (2011): "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences," *Journal of the European Economic Association*, 9(3), 522–550.

ECKEL, C. C., AND P. J. GROSSMAN (1998): "Are Women Less Selfish than Men?: Evidence from Dictator Experiments," *The Economic Journal*, 108(448), 726–735.

EHRHART, K.-M., AND M. OTT (2014): "Reference-Dependent Bidding in Dynamic Auctions," working paper.

EISENHUTH, R. (2010): "Auction Design with Loss-Averse Bidders: The Optimality of All-Pay Mechanisms," working paper.

ENDERS, W., AND G. A. HOOVER (2004): "Whose line is it? Plagiarism in economics," *Journal of Economic Literature*, 42(2), 487–493.

ERAT, S., AND U. GNEEZY (2012): "White lies," *Management Science*, 58(4), 723–733.

ERICSON, K. M. M., AND A. FUSTER (2011): "Expectations as Endowments: Evidence on Reference-Dependent Preferences from Exchange and Valuation Experiments," *Quarterly Journal of Economics*, 126, 1879–1907.

ERKAL, N., L. GANGADHARAN, AND N. NIKIFORAKIS (2011): "Relative Earnings and Giving in a Real-Effort Experiment," *American Economic Review*, 101(7), 3330–3348.

EUROPEAN COMMISSION (2004): "2004/138/EC: Commission Decision: Austrian Banks - Lombard Club," `http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:%2032004D0138:EN:NOT`.

FALK, A., E. FEHR, AND D. HUFFMAN (2008): "The power and limits of tournament incentives," working paper.

FISCHBACHER, U. (2007): "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments," *Experimental Economics*, 10, 171–178.

FUDENBERG, D., AND E. MASKIN (1986): "The Folk Theorem for Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533–554.

GABAIX, X., AND D. LAIBSON (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," *Quarterly Journal of Economics*, 121(2), 505–540.

GARICANO, L., AND I. PALACIOS-HUERTA (2006): "Sabotage in Tournaments: Making the Beautiful Game a Bit Less Beautiful," working paper.

GARRATT, R. J., C. WEINBERGER, AND N. JOHNSON (2013): "The State Street Mile: Age and Gender Differences in Competition Aversion in the Field," *Economic Inquiry*, 51(1), 806–815.

GILL, D., AND V. PROWSE (2012): "A Structural Analysis of Disappointment Aversion in a Real Effort Competition," *American Economic Review*, 102, 469–503.

GILL, D., AND R. STONE (2010): "Fairness and Desert in Tournaments," *Games and Economic Behavior*, 69, 346–364.

——— (2015): "Desert and inequity aversion in teams," *Journal of Public Economics*, 123, 42 – 54.

GNEEZY, U., K. L. LEONARD, AND J. A. LIST (2009): "Gender differences in competition: Evidence from a matrilineal and a patriarchal society," *Econometrica*, 77(5), 1637–1664.

GNEEZY, U., M. NIEDERLE, AND A. RUSTICHINI (2003): "Performance in Competitive Environments: Gender Differences," *The Quarterly Journal of Economics*, 118(3), 1049–1074.

GNEEZY, U., AND A. RUSTICHINI (2004): "Gender and Competition at a Young Age," *The American Economic Review*, 94(2), 377–381.

GREINER, B. (2004): "An Online Recruitment System for Economic Experiments," *Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht*, 63, 79–93.

GROSSMAN, S. (1981): "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24(3), 461.

GRUND, C., AND D. SLIWKA (2005a): "Envy and Compassion in Tournaments," *Journal of Economics & Management Strategy*, 14(1), 187–207.

GRUND, C., AND D. SLIWKA (2005b): "Envy and Compassion in Tournaments," *Journal of Economics and Management Strategy*, 14, 187–207.

GUL, F. (1991): "A Theory of Disappointment Aversion," *Econometrica*, 59, 667–686.

GÜNTHER, C., N. A. EKINCI, C. SCHWIEREN, AND M. STROBEL (2010): "Women Can't Jump? An Experiment on Competitive Attitudes and Stereotype Threat," *Journal of Economic Behavior & Organization*, 75(3), 395 – 401.

GÜRTLER, O. (2008): "On Sabotage in Collective Tournaments," *Journal of Mathematical Economics*, 44(3), 383–393.

GÜRTLER, O., J. MÜNSTER, AND P. NIEKEN (2013): "Information policy in tournaments with sabotage," *The Scandinavian Journal of Economics*, 115(3), 932–966.

GÜTH, W., C. SCHMIDT, AND M. SUTTER (2007): "Bargaining Outside the Lab - a Newspaper Experiment of a Three-Person Ultimatum Game," *The Economic Journal*, 117(518), 449–469.

HARBRING, C., AND B. IRLENBUSCH (2008): "How Many Winners are Good to Have?: On Tournaments with Sabotage," *Journal of Economic Behavior & Organization*, 65(3), 682 – 702.

——— (2011): "Sabotage in Tournaments: Evidence from a Laboratory Experiment," *Management Science*, 57(4), 611–627.

HARBRING, C., B. IRLENBUSCH, M. KRÄKEL, AND R. SELTEN (2007): "Sabotage in Corporate Contests - An Experimental Analysis," *International Journal of the Economics of Business*, 14(3), 367–392.

HART, O., AND B. HOLMSTRÖM (1987): "The Theory of Contracts," in *Advances in Economic Theory Fifth World Congress*, ed. by T. F. Bewlwy, pp. 70–155, New York. Cambridge University Press.

HEALY, A., AND J. PATE (2011): "Can Teams Help to Close the Gender Competition Gap?," *The Economic Journal*, 121(555), 1192–1204.

HEIDHUES, P., AND B. KŐSZEGI (2008): "Competition and Price Variation when Consumers are Loss Averse," *American Economic Review*, 98, 1245–1268.

——— (2014): "Regular Prices and Sales," *Theoretical Economics*, 9, 217–251.

HERWEG, F. (2013): "The Expectation-Based Loss-Averse Newsvendor," *Economics Letters*, 120, 429–432.

HERWEG, F., H. KARLE, AND D. MÜLLER (2014): "Incomplete Contracting, Renegotiation, and Expectation-Based Loss Aversion," working paper.

HERWEG, F., AND K. MIERENDORFF (2013): "Uncertain Demand, Consumer Loss Aversion, and Flat-Rate Tariffs," *Journal of the European Economic Association*, 11, 399–432.

HERWEG, F., D. MÜLLER, AND P. WEINSCHENK (2010): "Binary Payment Schemes: Moral Hazard and Loss Aversion," *American Economic Review*, 100, 2451–2477.

HOLT, C. A., AND S. K. LAURY (2002): "Risk Aversion and Incentive Effects," *The American Economic Review*, 92(5), 1644–1655.

IRIBERRI, N., AND P. REY-BIEL (2012): "Let's (Not) Talk about Sex: The Effect of Information Provision on Gender Differences in Performance under Competition," working paper.

KAHNEMAN, D., AND A. TVERSKY (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–292.

KARLE, H. (2014): "Creating Attachment through Advertising: Loss Aversion and Pre-Purchase Information," working paper.

KARLE, H., AND M. PEITZ (2014): "Competition under Consumer Loss Aversion," *RAND Journal of Economics*, 45, 1–31.

KŐSZEGI, B. (2010): "Utility from Anticipation and Personal Equilibrium," *Economic Theory*, 44, 415–444.

KŐSZEGI, B., AND M. RABIN (2006): "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 121, 1133–1165.

——— (2007): "Reference-Dependent Risk Attitudes," *American Economic Review*, 97, 1047–1073.

——— (2009): "Reference-Dependent Consumption Plans," *American Economic Review*, 99, 909–936.

KOSFELD, M., AND U. SCHÜWER (2011): “Add-on Pricing, Consumer Myopia and Regulatory Intervention,” Working Paper.

KRÄKEL, M. (2005): “Helping and Sabotaging in Tournaments,” *International Game Theory Review*, 7(02), 211–228.

——— (2008): “Emotions in Tournaments,” *Journal of Economic Behavior and Organization*, 67, 204–214.

KRÄKEL, M., AND P. NIEKEN (2012): “Relative Performance Pay in the Shadow of Crisis,” Working Paper.

KRÄKEL, M., AND P. NIEKEN (2015): “Relative Performance Pay in the Shadow of Crisis,” *European Economic Review*, 74, 244–268.

LAL, R., AND C. MATUTES (1994): “Retail Pricing and Advertising Strategies,” *Journal of Business*, 67(3), 345.

LANGE, A., AND A. RATAN (2010): “Multi-Dimensional Reference-Dependent Preferences in Sealed-Bid Auctions - How (most) Laboratory Experiments Differ from the Field,” *Games and Economic Behavior*, 68, 634–645.

LAZEAR, E., AND S. ROSEN (1981): “Rank-Order Tournaments as Optimum Labor Contracts,” *Journal of Political Economy*, 89, 841–864.

LAZEAR, E. P. (1989): “Pay Equality and Industrial Politics,” *Journal of Political Economy*, 97(3), 561–580.

LIST, J. A., C. D. BAILEY, P. J. EUZENT, AND T. L. MARTIN (2001): “Academic economists behaving badly? A survey on three areas of unethical behavior,” *Economic Inquiry*, 39(1), 162–170.

LOOMES, G., AND R. SUGDEN (1986): “Disappointment and Dynamic Consistency in Choice under Uncertainty,” *Review of Economic Studies*, 53, 271–282.

LUDWIG, S., P. WICHARDT, AND H. WICKHORST (2011): “Overconfidence can Improve an Agent’s Relative and Absolute Performance,” *Economics Letters*, 110, 193–196.

LUNDEBERG, M. A., P. W. FOX, AND J. PUNĆCOHAŔ (1994): “Highly confident but wrong: gender differences and similarities in confidence judgments.,” *Journal of educational psychology*, 86(1), 114.

MACERA, R. (2013): “Intertemporal Incentives With Expectation-Based Reference-Dependent Preferences,” working paper.

MAILATH, G., AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, USA.

MANNING, W. G., N. DUAN, AND W. H. ROGERS (1987): “Monte Carlo evidence on the choice between sample selection and two-part models,” *Journal of econometrics*, 35(1), 59–82.

MASCLET, D., E. PETERLE, AND S. LARRIBEAU (2012): “Gender Differences in Competitive and Non Competitive Environments: An Experimental Evidence,” working paper.

MIAO, C. (2010): “Consumer myopia, standardization and aftermarket monopolization,” *European Economic Review*, 54, 931–946.

MILGROM, P. (1981): “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 12(2), 380–391.

MIRAVETE, E. (2007): “The Doubtful Profitability of Foggy Pricing,” Working Paper.

MÜNSTER, J. (2007): “Selection Tournaments, Sabotage, and Participation,” *Journal of Economics & Management Strategy*, 16(4), 943–970.

NIEDERLE, M., AND L. VESTERLUND (2007): “Do Women Shy Away from Competition? Do Men Compete Too Much?,” *The Quarterly Journal of Economics*, 122(3), 1067–1101.

——— (2011): “Gender and Competition,” *Annual Review of Economics*, 3(1), 601–630.

O’DONOGHUE, T., AND M. RABIN (1999): “Doing It Now or Later,” *American Economic Review*, 89(1), 103–124.

PRESTON, I., AND S. SZYMANSKI (2003): “Cheating in Contests,” *Oxford Review of Economic Policy*, 19(4), 612–624.

PRICE, C. R. (2012): “Gender, Competition, and Managerial Decisions,” *Management Science*, 58(1), 114–122.

PRICE, J. (2008): “Gender Differences in the Response to Competition,” *Industrial and Labor Relations Review*, 61(3), 320–333.

REUBEN, E., P. REY-BIEL, P. SAPIENZA, AND L. ZINGALES (2012): “The emergence of male leadership in competitive environments,” *Journal of Economic Behavior & Organization*, 83(1), 111–117.

ROSATO, A. (2014a): “Loss Aversion in Sequential Auctions: Endogenous Interdependence, Informational Externalities and the “Afternoon Effect“,” working paper.

——— (2014b): "Selling Substitute Goods to Loss-Averse Consumers: Limited Availability, Bargains and Rip-offs," working paper.

ROSEN, S. (1986): "Prizes and Incentives in Elimination Tournaments," *American Economic Review*, 76, 701–715.

ROSENTHAL, R. W. (1979): "Sequences of Games with Varying Opponents," *Econometrica*, 47, 1353–1366.

RUBINSTEIN, A. (1991): "Comments on the Interpretation of Game Theory," *Econometrica*, 59, 909–924.

SANTOS-PINTO, L. (2010): "Positive Self-Image in Tournaments," *International Economic Review*, 51, 475–496.

SCHWARTZ, B. (2000): "Self-Determination: The Tyranny of Freedom," *American Psychologist*, 55(1), 79–88.

SCHWARTZ, S. H., G. MELECH, A. LEHMANN, S. BURGESS, M. HARRIS, AND V. OWENS (2001): "Extending the cross-cultural validity of the theory of basic human values with a different method of measurement," *Journal of cross-cultural Psychology*, 32(5), 519–542.

SCHWARTZ, S. H., AND T. RUBEL (2005): "Sex Differences in Value Priorities: Cross-Cultural and Multimethod Studies," *Journal of Personality and Social Psychology*, 89(6), 1010–1028.

SCHWIEREN, C., AND D. WEICHSELBAUMER (2010): "Does Competition Enhance Performance or Cheating? A Laboratory Experiment," *Journal of Economic Psychology*, 31(3), 241–253.

SHALEV, J. (2000): "Loss Aversion Equilibrium," *International Journal of Game Theory*, 29, 269–287.

SHAPIRO, C. (1995): "Aftermarkets and Consumer Welfare: Making Sense of Kodak," *Antitrust Law Journal*, 63, 483–511.

SHLEIFER, A. (2004): "Does Competition Destroy Ethical Behavior?," *American Economic Review*, 94(2), 414–418.

TNS (2011): "Special Eurobarometer 342: Consumer empowerment," TNS Opinion & Social.

VAN DIJK, F., J. SONNEMANS, AND F. VAN WINDEN (2001): "Incentive Systems in a Real Effort Experiment," *European Economic Review*, 45(2), 187–214.

VANDEGRIFT, D., AND P. BROWN (2005): "Gender Differences in the Use of High-Variance Strategies in Tournament Competition," *The Journal of Socio-Economics*, 34(6), 834 – 849.

VANDEGRIFT, D., AND A. YAVAS (2010): "An Experimental Test of Sabotage in Tournaments," *Journal of Institutional and Theoretical Economics JITE*, 166(2), 259–285.

WIELAND, A., AND R. SARIN (2012): "Domain Specificity of Sex Differences in Competition," *Journal of Economic Behavior & Organization*, 83(1), 151–157.

WOZNIAK, D., W. T. HARBAUGH, AND U. MAYR (2010): "Choices about Competition: Differences by Gender and Hormonal Fluctuations, and the Role of Relative Performance Feedback," working paper.