

Dato, Simon; Grunewald, Andreas; Müller, Daniel

Working Paper

Expectation-Based Loss Aversion and Strategic Interaction

Bonn Econ Discussion Papers, No. 02/2015

Provided in Cooperation with:

Bonn Graduate School of Economics (BGSE), University of Bonn

Suggested Citation: Dato, Simon; Grunewald, Andreas; Müller, Daniel (2015) : Expectation-Based Loss Aversion and Strategic Interaction, Bonn Econ Discussion Papers, No. 02/2015, University of Bonn, Bonn Graduate School of Economics (BGSE), Bonn

This Version is available at:

<https://hdl.handle.net/10419/128617>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Bonn Econ Discussion Papers

Discussion Paper 02/2015

Expectation-Based Loss Aversion and Strategic Interaction

by

Simon Dato, Andreas Grunewald and Daniel Müller

February 2015



Bonn Graduate School of Economics
Department of Economics
University of Bonn
Kaiserstrasse 1
D-53113 Bonn

Financial support by the
Deutsche Forschungsgemeinschaft (DFG)
through the
Bonn Graduate School of Economics (BGSE)
is gratefully acknowledged.

Expectation-Based Loss Aversion and Strategic Interaction*

SIMON DATO[†], ANDREAS GRUNEWALD[‡], DANIEL MÜLLER[§]

February 2, 2015

This paper 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. Third, we analyze equilibrium play under expectation-based loss aversion and comment on the existence of equilibria.

JEL classification: C72, D01, D03, D81

Keywords: Non-Cooperative Games, Expectation-Based Loss Aversion, Reference-Dependent Preferences, Mixed Strategies

*We thank Armin Falk, Paul Heidhues, Fabian Herweg, Daniel Krähmer, Matthias Kräkel, Simone Quercia, and Matthew Rabin for helpful comments and suggestions. Furthermore, we have benefited from comments made by the seminar audience at the University of Bonn. All errors are of course our own.

[†]University of Bonn, Institute for Applied Microeconomics, Adenauerallee 24-42, D-53113 Bonn, Germany, E-mail address: simdato@uni-bonn.de, Tel: +49-228-739214

[‡]University of Bonn, Institute for Applied Microeconomics, Adenauerallee 24-42, D-53113 Bonn, Germany, E-mail address: gruni@uni-bonn.de, Tel: +49-228-739217.

[§]University of Bonn, Institute for Applied Microeconomics, Adenauerallee 24-42, D-53113 Bonn, Germany, E-mail address: daniel.mueller@uni-bonn.de, Tel: +49-228-739212, Fax: +49-228-739212.

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 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.

¹The 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.

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

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

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

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, 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 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 formally introduces the class of games we study while Section 4 extends the equilibrium concepts PE and CPE to strategic interaction. In Section 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 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 7. Section 8 concludes.

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

⁴As 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.

⁵Karle 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.

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 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

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

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

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 resembling those of players with standard preferences, equilibrium existence in Shalev (2000) is guaranteed by Kakutani's fixed point theorem.

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 $\mathcal{I} = \{1, \dots, I\}$ is finite. Second, each player $i \in \mathcal{I}$ has a finite pure-strategy space $\mathcal{S}^i = \{s_1^i, \dots, s_{M^i}^i\}$. A pure-strategy profile is denoted by $s = (s^1, \dots, s^I) \in \mathcal{S}$, where $\mathcal{S} = \times_{i=1}^I \mathcal{S}^i$. Third, there is a finite set $\Theta = \{\theta_1, \dots, \theta_N\}$, 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 $u^i : \mathcal{S} \times \Theta \rightarrow \mathcal{U}^i \subset \mathbb{R}$ 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 $\sigma^i = (\sigma^i(s_1^i), \dots, \sigma^i(s_{M^i}^i))$ for player $i \in \mathcal{I}$ is a lottery over her pure strategies, where $\sigma^i(s_m^i)$ denotes the probability of player i playing the pure strategy s_m^i . The space of player i 's mixed strategies is denoted by Σ^i . Accordingly, the space of mixed-strategy profiles $\sigma = (\sigma^1, \dots, \sigma^I)$ 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

⁸Clearly, 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}) .

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:

$$\mu(u - u^r) = \begin{cases} \eta(u - u^r) & \text{if } u \geq u^r \\ \eta\lambda(u - u^r) & \text{if } u < u^r \end{cases}. \quad (1)$$

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 $\Lambda^i(u) = \{(s, \theta) \in S \times \Theta \mid u^i(s, \theta) = u\}$ 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 $P^i(u|\sigma) = \sum_{(s, \theta) \in \Lambda^i(u)} Q(\theta) \prod_{j=1}^I \sigma^j(s^j)$. 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

$$U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i}) = \sum_{u \in \mathcal{U}^i} P^i(u|(\sigma^i, \sigma^{-i})) \cdot u + \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|(\sigma^i, \sigma^{-i})) \cdot P^i(\tilde{u}|(\hat{\sigma}^i, \sigma^{-i})) \cdot \mu(u - \tilde{u}). \quad (2)$$

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 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

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.

⁹Almost all of the contributions cited in Section 2 use this piece-wise linear specification of the gain-loss function.

¹⁰With 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).

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 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 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 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 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 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 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




Figure 1: *Material utility payoff matrix (left panel) and, given standard expected utility preferences, player 1's best response curve (right panel) in the Chicken game.*

the other players' strategies. Hence, in a CPNE, every player plays a CPE in the sense of Kőszegi and Rabin (2007).¹¹

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 1.

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 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 $\mathcal{S}^i = \{\text{go straight, swerve}\}$. To ease notation, we denote $\sigma^1(\text{go straight}) = \alpha_1$ and $\sigma^2(\text{go straight}) = \beta_1$, where $0 \leq \alpha_1, \beta_1 \leq$

¹¹According Definitions 1 and 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 5 remain valid under such a modification.

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 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.

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

$$L^i(\sigma^i, \sigma^{-i}) = (P(u|\sigma^i, \sigma^{-i}))_{u \in \mathcal{U}^i} \quad (3)$$

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

A pure strategy s_k^i is redundant if the lottery over material utility outcomes induced by s_k^i is a linear combination of the lotteries induced by a set of pure strategies not containing s_k^i . Note that pure strategy s_k^i being non-redundant implies that player i cannot replicate the lottery over outcomes induced by s_k^i by playing any mixed strategy excluding s_k^i . In this sense, only a nonredundant 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 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*

¹²To 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.

under L^j . Then

$$\frac{\sum_{u \in \mathcal{U}} \sum_{\tilde{u} \in \mathcal{U}} p^A(u) p^A(\tilde{u}) |u - \tilde{u}| + \sum_{u \in \mathcal{U}} \sum_{\tilde{u} \in \mathcal{U}} p^B(u) p^B(\tilde{u}) |u - \tilde{u}|}{2} \leq \sum_{u \in \mathcal{U}} \sum_{\tilde{u} \in \mathcal{U}} p^A(u) p^B(\tilde{u}) |u - \tilde{u}|, \quad (4)$$

with (4) holding with equality if and only if L^A and L^B are identical.

Lemma 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 1 relates to the above definition of redundancy: If two pure strategies s_k^i and s_m^i induce lotteries $L^i(s_k^i, \sigma^{-i})$ and $L^i(s_m^i, \sigma^{-i})$ for which (4) holds with equality, these lotteries are identical and the pure strategies s_k^i and s_m^i are redundant.

Lemma 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 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 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.

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

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 $\sigma^i(s_m^i)$ of player i 's strategy σ^i . Hence, the marginal utility of $U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i})$ with respect to $\sigma^i(s_m^i)$ 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_m^i to some other pure strategy s_k^i changes expected utility at a constant rate equal to $\partial U^i / \partial \sigma^i(s_k^i) - \partial U^i / \partial \sigma^i(s_m^i)$. 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_m^i is a PE for player i if and only if the marginal expected utility with respect to component $\sigma^i(s_m^i)$ is among the greatest marginal expected utilities given $\hat{\sigma}^i(s_m^i) = 1$ and $\hat{\sigma}^i(s_k^i) = 0$ for all $k \neq m$. Hence, s_m^i is at least as attractive as any other pure strategy since moving probability mass from s_m^i 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

$$\Gamma(\sigma^i) = \{s_m^i \in \mathcal{S}^i \mid \sigma^i(s_m^i) > 0\} \quad (5)$$

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

$$R^i(\sigma^{-i}) = \{\sigma^i \in \Sigma^i \mid U^i(\sigma^i, \sigma^i, \sigma^{-i}) \geq U^i(\tilde{\sigma}^i, \sigma^i, \sigma^{-i}), \forall \tilde{\sigma}^i \in \Sigma^i\} \quad (6)$$

the correspondence of PEs for player i . Proposition 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

¹⁴Proposition 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.

Proposition 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 $|\Gamma(\bar{\sigma}^i)| \geq 2$, 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 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 1(i), if a mixed strategy PE involves player i mixing only over nonredundant 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.

Second, with regard to Proposition 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 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 (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

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

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

$$\frac{\partial U^1(\sigma^1, \hat{\sigma}^1, \sigma^2)}{\partial \alpha_1} = 3\beta_2 + \eta \left\{ \beta_2^2 \hat{\alpha}_2 + \beta_1\beta_2(1 - \lambda)[3\hat{\alpha}_1 + 2\hat{\alpha}_2] - \lambda\hat{\alpha}_2\beta_1^2 \right\} \quad (7)$$

$$\frac{\partial U^1(\sigma^1, \hat{\sigma}^1, \sigma^2)}{\partial \alpha_2} = 1 + \beta_2 + \eta \left\{ \hat{\alpha}_1\beta_1^2 + \beta_1\beta_2[(1 - \lambda)\hat{\alpha}_2 + 2\hat{\alpha}_1] + \beta_2^2(1 - \lambda)\hat{\alpha}_1 \right\} \quad (8)$$

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 (7) is larger than the one in (8) 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 (7) is smaller than the one in (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 (7) and (8) have to be identical for expectations $\hat{\alpha}_1 = \alpha_1$ and $\hat{\alpha}_2 = \alpha_2$.

The left panel of Figure 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




Figure 2: *player 1's set of PEs (left panel) and her set of CPEs (right panel) in the Chicken game for $\lambda = 3$ and $\eta = 1$.*

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 1 also become apparent in Figure 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 $\frac{d^2 U^1(\sigma^1, \hat{\sigma}^1, \sigma^2)}{d\alpha_j d\hat{\alpha}_j} \geq 0$, $j \in \{1, 2\}$. 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 $\beta_1 \in (\underline{\beta}, \frac{1}{2}]$ 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.

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 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 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

$$U^1(\sigma^1, \sigma^1, \sigma^2) = 3\alpha_1\beta_2 + \alpha_2\beta_1 + 2\alpha_2\beta_2 \\ - \eta(\lambda - 1) \left[3\alpha_1^2\beta_1\beta_2 + 2(2\alpha_1\alpha_2\beta_1\beta_2) + \alpha_1\alpha_2\beta_1^2 + \alpha_1\alpha_2\beta_2^2 + \alpha_1^2\beta_1\beta_2 \right].$$

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 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.

6. EQUILIBRIUM EXISTENCE AND BEHAVIOR

Section 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 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 2), it follows that the set of pure strategy CPNEs is also identical to the set of pure strategy Nash equilibria.

Proposition 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.

6.1. Personal Nash Equilibrium

As we have seen in Section 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 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

¹⁵This 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.

\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 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 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 5. The middle panel of Figure 3 depicts the sets of PEs for both players. According to Definition 1, the game's PNEs lie at the intersections of the two PE correspondences. As implied by Proposition 3, the Nash equilibria in 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

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




Figure 3: The left panel shows the best responses for players 1 and 2 with standard preferences in the Chicken game depicted in Figure 1, the middle one their sets of PEs, and the right panel their sets of CPNEs, if $\lambda = 3$ and $\eta = 1$. The corresponding equilibria are at the intersection of the best responses.

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.

6.2. Choice-Acclimating Personal Nash Equilibrium

In Proposition 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 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., $L^i(s_k^i, \sigma^{-i}) \neq L^i(s_m^i, \sigma^{-i})$ for all $i \in \mathcal{I}$, $\sigma^{-i} \in \Sigma^{-i}$, and $s_k^i, s_m^i \in \mathcal{S}^i$ with $s_k^i \neq s_m^i$. 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 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 (2014). 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 1(ii), which follows from Propositions 2 and 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 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 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 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 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.*

Illustration: Equilibrium Behavior for Choice Acclimating Beliefs

The right panel of Figure 3 shows the set of CPEs of the two drivers in the Chicken game. As implied by Proposition 3, the two pure strategy Nash equilibria of the game also constitute CPNEs. Moreover, Figure 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

¹⁷To be precise, most of the above applications comprise multidimensional outcomes. In Proposition 6, we show that our results carry over to the case of multidimensional outcomes.

with standard preferences or to situations with fixed expectations, there exist only two CPNEs in the Chicken game.

7. DISCUSSION

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

¹⁸For an example how these interpretations generate additional equilibria reconsider the large population

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.*

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 $u^i : \mathcal{S} \times \Theta \rightarrow \mathcal{U}^i \subset \mathbb{R}^R$ 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, $u^i(s, \theta) = (u_1^i(s, \theta), \dots, u_R^i(s, \theta)) \in \mathbb{R}^R$. $P^i(u|\sigma)$ 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

$$U^i(\sigma^i, \hat{\sigma}^i, \sigma^{-i}) = \sum_{u \in \mathcal{U}^i} P^i(u | (\sigma^i, \sigma^{-i})) \cdot \sum_{r=1}^R u_r \\ + \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u | (\sigma^i, \sigma^{-i})) \cdot P^i(\tilde{u} | (\hat{\sigma}^i, \sigma^{-i})) \cdot \sum_{r=1}^R \mu(u_r - \tilde{u}_r). \quad (9)$$

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,

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 $(\alpha_1, \alpha_2) = (\beta_1, \beta_2) = (0.19, 0.81)$.

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

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 5 and 6 carry over to the case of multidimensional payoffs.

Proposition 6. *Suppose material payoffs are multidimensional. Then the results from Proposition 1, Proposition 2, Proposition 4, Corollary 1 and Proposition 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 shows, more outcomes can be supported in equilibrium.

Proposition 7. *Suppose there is no draw of nature, $\Theta = \{\tilde{\theta}\}$, and all players' payoffs are multidimensional, $U^i \subset \mathbb{R}^R$ 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 $r^i(\tilde{s}^i) = 1, \dots, R$ such that $u_{r^i(\tilde{s}^i)}^i((s^i, s^{-i}), \tilde{\theta}) > u_{r^i(\tilde{s}^i)}^i((\tilde{s}^i, s^{-i}), \tilde{\theta})$.*

Without inherent uncertainty in the game, the logic underlying parts (i) and (ii) of Proposition 7 is the same as for the corresponding statements regarding one-dimensional payoffs in Proposition 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 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 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.

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 multi-dimensional. 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 4.

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

²⁰As an example consider a slightly asymmetric matching pennies game, in which the only PNE involves mixed strategies. As shown in the proof of Proposition 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.

mixed strategies would allow a more comprehensive study of equilibria in these contexts.

REFERENCES

- ABELER, J., A. FALK, L. GÖTTE, AND D. HUFFMAN (2011): “Reference Points and Effort Provision,” *American Economic Review*, 101, 470–492.
- 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.
- 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.
- 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.
- 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.
- DATO, S., A. GRUNEWALD, AND D. MÜLLER (2014): “Expectation-Based Loss Aversion and Rank-Order Tournaments,” working paper.
- 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.
- 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.
- 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.
- GUL, F. (1991): “A Theory of Disappointment Aversion,” *Econometrica*, 59, 667–686.
- 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.
- 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.
- 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.
- LOOMES, G., AND R. SUGDEN (1986): “Disappointment and Dynamic Consistency in Choice under Uncertainty,” *Review of Economic Studies*, 53, 271–282.
- MACERA, R. (2013): “Intertemporal Incentives With Expectation-Based Reference-Dependent Preferences,” working paper.
- 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.
- 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.
- SHALEV, J. (2000): “Loss Aversion Equilibrium,” *International Journal of Game Theory*, 29, 269–287.

APPENDIX

A. PROOFS OF PROPOSITIONS AND LEMMAS

Proof of Lemma 1. Suppose \mathcal{U} is an outcome space with $N \geq 2$ elements, $u_1 > u_2 > \dots > u_N$. Furthermore, let $L^A = \{p_1^A, p_2^A, \dots, p_N^A\}$ and $L^B = \{p_1^B, p_2^B, \dots, p_N^B\}$ 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 $k = 1, \dots, N$, $\mathcal{U}(k) = \{u_1, \dots, u_k\}$ denotes the “truncated” outcome space which contains only the k highest elements of \mathcal{U} . For $k = 1, \dots, N$ and $j \in \{A, B\}$, $L^j(k) = (p_1^j(k), \dots, p_k^j(k))$ 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

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

with respect to u_k yields

$$\frac{df(L^A(k), L^B(k))}{du_k} = [p_k^A(k) - p_k^B(k)]^2 \geq 0. \quad (\text{A.1})$$

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.1),

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

where (A.2) 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 $k = 1, \dots, N$. Given that $p_t^A(t) = p_t^B(t)$ for all $t = k+1, \dots, N$, 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 $k = 1, \dots, N$. Conversely, if $p_k^A \neq p_k^B$ for some $k = 1, \dots, N$, then $f(L^A, L^B) < 0$. \square

Proof of Proposition 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 $\hat{\sigma}^i$:

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

with

$$q_{km}^i(\sigma^{-i}) = \sum_{\theta \in \Theta} Q(\theta) \left\{ \sum_{s^{-i} \in \mathcal{S}^{-i}} \left(\prod_{(s^j)_{j \neq i} \in \mathcal{S}^{-i}} \sigma^j(s^j) \right) \left[\sum_{\tilde{\theta} \in \Theta} Q(\tilde{\theta}) \left(\sum_{\tilde{s}^{-i} \in \mathcal{S}^{-i}} \left(\prod_{j \neq i} \sigma^j(\tilde{s}^j) \right) \mu(u((s_k^i, s^{-i}), \theta) - u((s_m^i, \tilde{s}^{-i}), \tilde{\theta})) \right) \right] \right\}$$

and

$$b_k^i(\sigma^{-i}) = \sum_{\theta \in \Theta} Q(\theta) \left[\sum_{s^{-i} \in \mathcal{S}^{-i}} \left(\prod_{j \neq i} \sigma^j(s^j) \right) u((s_k^i, s^{-i}), \theta) \right],$$

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

$$\begin{pmatrix} MU_1^i(\hat{\sigma}^i, \sigma^{-i}) \\ \vdots \\ MU_{M^i}^i(\hat{\sigma}^i, \sigma^{-i}) \end{pmatrix} = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \dots & a_{1M^i}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{M^i1}(\sigma^{-i}) & \dots & a_{M^iM^i}(\sigma^{-i}) \end{pmatrix}}_{=: A(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s_1^i) \\ \vdots \\ \hat{\sigma}^i(s_{M^i}^i) \end{pmatrix}, \quad (\text{A.3})$$

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., $\bar{\sigma}^i(s_k^i) > 0$ for $k = 1, \dots, m$ and $\bar{\sigma}^i(s_k^i) = 0$ for $k > m$, where $\sum_{k=1}^m \bar{\sigma}^i(s_k^i) = 1$. As described in the text, $MU_k^i(\hat{\sigma}^i, \sigma^{-i})$ reflects the attractiveness to play pure strategy s_k^i . Since $\bar{\sigma}^i$ is assumed to be a PE, $MU_1^i(\bar{\sigma}^i, \sigma^{-i}) = \dots = MU_m^i(\bar{\sigma}^i, \sigma^{-i}) = \bar{u} \geq \max_{k > m} MU_k^i(\bar{\sigma}^i, \sigma^{-i})$. With $\bar{\sigma}^i(s_k^i) = 0$ for $k > m$, the mixing probabilities $\bar{\sigma}^i(s_1^i), \dots, \bar{\sigma}^i(s_m^i)$ are thus a solution of the following system of linear equations:

$$\begin{pmatrix} \bar{u} \\ \vdots \\ \bar{u} \end{pmatrix} = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \dots & a_{1m}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{m1}(\sigma^{-i}) & \dots & a_{mm}(\sigma^{-i}) \end{pmatrix}}_{=: A'(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s_1^i) \\ \vdots \\ \hat{\sigma}^i(s_m^i) \end{pmatrix}. \quad (\text{A.4})$$

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.4) has a unique solution, which (by hypothesis) is given by the vector $(\bar{\sigma}^i(s_1^i), \dots, \bar{\sigma}^i(s_m^i))$

with $\sum_{k=1}^m \sigma^i(s_k^i) = 1$. In contradiction to the statement, suppose that there exists a different PE, $(\tilde{\sigma}^i(s_1^i), \dots, \tilde{\sigma}^i(s_{M^i}^i)) \neq (\bar{\sigma}^i(s_1^i), \dots, \bar{\sigma}^i(s_{M^i}^i))$, that mixes over the same set of pure strategies, i.e., $\tilde{\sigma}^i(s_k^i) > 0$ for $k = 1, \dots, m$ 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 $(\tilde{\sigma}^i(s_1^i), \dots, \tilde{\sigma}^i(s_m^i))$ solves a system of linear equations

$$\begin{pmatrix} \tilde{u} \\ \vdots \\ \tilde{u} \end{pmatrix} = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \dots & a_{1m}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{m1}(\sigma^{-i}) & \dots & a_{mm}(\sigma^{-i}) \end{pmatrix}}_{=A'(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s_1^i) \\ \vdots \\ \hat{\sigma}^i(s_m^i) \end{pmatrix}. \quad (\text{A.5})$$

By full rank of $A'(\sigma^{-i})$, we must have $\tilde{u} \neq \bar{u}$, because otherwise $(\tilde{\sigma}^i(s_1^i), \dots, \tilde{\sigma}^i(s_m^i)) = (\bar{\sigma}^i(s_1^i), \dots, \bar{\sigma}^i(s_m^i))$. Hence either \bar{u} or \tilde{u} differs from zero. Suppose, that $\tilde{u} \neq 0$. In consequence, (A.5) implies that $(\bar{\sigma}^i(s_1^i), \dots, \bar{\sigma}^i(s_m^i)) = (\frac{\bar{u}}{\tilde{u}}\tilde{\sigma}^i(s_1^i), \dots, \frac{\bar{u}}{\tilde{u}}\tilde{\sigma}^i(s_m^i))$. 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 $a_k(\sigma^{-i}) = (a_{k1}(\sigma^{-i}), \dots, a_{km}(\sigma^{-i}))$. 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 $a_1(\sigma^{-i}), \dots, a_{m-1}(\sigma^{-i})$, i.e., $a_m(\sigma^{-i}) = \sum_{k=1}^{m-1} \gamma_k a_k(\sigma^{-i})$ for some numbers $\gamma_1, \dots, \gamma_{m-1}$. Therefore

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

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 $k = 1, \dots, m$, 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

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

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

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

where

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

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 $k = 1, \dots, m$ and $\sigma_{\tilde{x}}^i(s_k^i) = 0$ for at least one $k = 1, \dots, m$. 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

$$E[L^i(\bar{\sigma}^i, \sigma^{-i})] - E[L^i(\sigma_{\tilde{x}}^i, \sigma^{-i})] = GL^i(\sigma_{\tilde{x}}^i, \sigma_{\tilde{x}}^i, \sigma^{-i}) - GL^i(\bar{\sigma}^i, \sigma_{\tilde{x}}^i, \sigma^{-i}). \quad (\text{A.6})$$

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

$$E[L^i(\bar{\sigma}^i, \sigma^{-i})] - E[L^i(\sigma_{\tilde{x}}^i, \sigma^{-i})] = GL^i(\sigma_{\tilde{x}}^i, \bar{\sigma}^i, \sigma^{-i}) - GL^i(\bar{\sigma}^i, \bar{\sigma}^i, \sigma^{-i}). \quad (\text{A.7})$$

(A.6) and (A.7) together imply

$$\begin{aligned} & GL^i(\sigma_{\tilde{x}}^i, \bar{\sigma}^i, \sigma^{-i}) - GL^i(\bar{\sigma}^i, \bar{\sigma}^i, \sigma^{-i}) - GL^i(\sigma_{\tilde{x}}^i, \sigma_{\tilde{x}}^i, \sigma^{-i}) + GL^i(\bar{\sigma}^i, \sigma_{\tilde{x}}^i, \sigma^{-i}) = 0 \quad (\text{A.8}) \\ & \Leftrightarrow \frac{1}{2} \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u | (\bar{\sigma}^i, \sigma^{-i})) \cdot P^i(\tilde{u} | (\bar{\sigma}^i, \sigma^{-i})) \cdot |u - \tilde{u}| \\ & - \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u | (\bar{\sigma}^i, \sigma^{-i})) \cdot P^i(\tilde{u} | (\sigma_{\tilde{x}}^i, \sigma^{-i})) \cdot |u - \tilde{u}| \\ & + \frac{1}{2} \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u | (\sigma_{\tilde{x}}^i, \sigma^{-i})) \cdot P^i(\tilde{u} | (\sigma_{\tilde{x}}^i, \sigma^{-i})) \cdot |u - \tilde{u}| = 0. \end{aligned}$$

By Lemma 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 σ_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 $\gamma(s_j^i) = \frac{\sigma_x^i(s_j^i) - \bar{\sigma}^i(s_j^i)}{\bar{\sigma}^i(s_m^i)}$ for $j = 1, \dots, m-1$ 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

$$z(\hat{\sigma}^i, \sigma^{-i}) = \underbrace{\begin{pmatrix} a_{11}(\sigma^{-i}) & \dots & a_{1m}(\sigma^{-i}) \\ \vdots & & \vdots \\ a_{m1}(\sigma^{-i}) & \dots & a_{mm}(\sigma^{-i}) \end{pmatrix}}_{=A'(\sigma^{-i})} \begin{pmatrix} \hat{\sigma}^i(s_1^i) \\ \vdots \\ \hat{\sigma}^i(s_m^i) \end{pmatrix} - \begin{pmatrix} \bar{u} \\ \vdots \\ \bar{u} \end{pmatrix}, \quad (\text{A.9})$$

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 $g : \Sigma^{-i} \rightarrow \mathbb{R}^m$ 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, $MU_1^i(\bar{\sigma}^i, \sigma^{-i}) = \dots = MU_m^i(\bar{\sigma}^i, \sigma^{-i}) > \max_{k>m} MU_k^i(\bar{\sigma}^i, \sigma^{-i})$, $\bar{\sigma}^i(s_k^i) > 0$ 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 $(\hat{\sigma}_\varepsilon^i(s_1^i), \dots, \hat{\sigma}_\varepsilon^i(s_m^i))$, are also strictly positive. Hence, for the mixed strategy $\sigma_\varepsilon^i = (\sigma_\varepsilon^i(s_1^i), \dots, \sigma_\varepsilon^i(s_{M^i}^i))$ with

$$\sigma_\varepsilon^i(s_k^i) = \begin{cases} \frac{\hat{\sigma}_\varepsilon^i(s_k^i)}{\sum_{j=1}^m \hat{\sigma}_\varepsilon^i(s_j^i)} & \text{if } k \leq m \\ 0 & \text{if } k > m \end{cases}, \quad (\text{A.10})$$

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

Proof of Proposition 2. Suppose $\bar{\sigma}^i \in \Sigma^i$ is a mixed CPE with $\bar{\sigma}^i(s_k^i) > 0$ for $k = 1, \dots, m$ (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

$$\sigma_{m'}(s_{m'}^i) = \bar{\sigma}^i(s_{m'}^i) + \bar{\sigma}^i(s_{m''}^i), \sigma_{m'}(s_{m''}^i) = 0, \sigma_{m'}(s_k^i) = \bar{\sigma}^i(s_k^i) \text{ for } k \neq m', m'' \quad (\text{A.11})$$

and

$$\sigma_{m''}(s_{m'}^i) = 0, \sigma_{m''}(s_{m''}^i) = \bar{\sigma}^i(s_{m'}^i) + \bar{\sigma}^i(s_{m''}^i), \sigma_{m''}(s_k^i) = \bar{\sigma}^i(s_k^i) \text{ for } k \neq m', m'', \quad (\text{A.12})$$

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

$$\begin{aligned} & \beta U^i(\sigma_{m'}, \sigma_{m'}, \sigma^{-i}) + (1-\beta)U^i(\sigma_{m''}, \sigma_{m''}, \sigma^{-i}) \\ & > U^i(\beta\sigma_{m'} + (1-\beta)\sigma_{m''}, \beta\sigma_{m'} + (1-\beta)\sigma_{m''}, \sigma^{-i}) \\ \Leftrightarrow & \beta GL^i(\sigma_{m'}, \sigma_{m'}, \sigma^{-i}) + (1-\beta)GL^i(\sigma_{m''}, \sigma_{m''}, \sigma^{-i}) \\ & > GL^i(\beta\sigma_{m'} + (1-\beta)\sigma_{m''}, \beta\sigma_{m'} + (1-\beta)\sigma_{m''}, \sigma^{-i}) \\ = & \beta^2 GL^i(\sigma_{m'}, \sigma_{m'}, \sigma^{-i}) + (1-\beta)^2 GL^i(\sigma_{m''}, \sigma_{m''}, \sigma^{-i}) \\ & + \beta(1-\beta)[GL^i(\sigma_{m'}, \sigma_{m''}, \sigma^{-i}) + GL^i(\sigma_{m''}, \sigma_{m'}, \sigma^{-i})] \\ \Leftrightarrow & -GL^i(\sigma_{m'}, \sigma_{m'}, \sigma^{-i}) + GL^i(\sigma_{m''}, \sigma_{m''}, \sigma^{-i}) \\ & + GL^i(\sigma_{m'}, \sigma_{m''}, \sigma^{-i}) - GL^i(\sigma_{m''}, \sigma_{m'}, \sigma^{-i}) < 0 \tag{A.13} \\ \Leftrightarrow & \frac{1}{2} \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|\sigma_{m'}, \sigma^{-i}) \cdot P^i(\tilde{u}|\sigma_{m'}, \sigma^{-i}) \cdot |u - \tilde{u}| \\ & - \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|\bar{\sigma}_{m'}^i, \sigma^{-i}) \cdot P^i(\tilde{u}|\sigma_{m''}^i, \sigma^{-i}) \cdot |u - \tilde{u}| \\ & + \frac{1}{2} \sum_{u \in \mathcal{U}^i} \sum_{\tilde{u} \in \mathcal{U}^i} P^i(u|\sigma_{m''}^i, \sigma^{-i}) \cdot P^i(\tilde{u}|\sigma_{m''}^i, \sigma^{-i}) \cdot |u - \tilde{u}| < 0. \end{aligned}$$

By Lemma 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

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

and

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

by (A.11) and (A.12) we have

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

Hence, for $\bar{\sigma}^i$ to be a CPE it must hold that $P^i(u|s_{m'}^i, \sigma^{-i}) = P^i(u|s_{m''}^i, \sigma^{-i})$ 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. \square

Proof of Proposition 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 $k = 1, \dots, M^i$, such that $U^i(s^i, s^i, s^{-i}) \geq U^i(s_k^i, s_k^i, s^{-i})$ for all $k = 1, \dots, M^i$. Together with Proposition 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 2.

Conversely, suppose the pure strategy profile (s^i, s^{-i}) constitutes a CPNE. By Definition 2, then $U^i(s^i, s^i, s^{-i}) \geq U^i(s_k^i, s_k^i, s^{-i})$ for all $k = 1, \dots, M^i$, which implies that $u^i((s^i, s^{-i}), \tilde{\theta}) \geq u^i((s_k^i, s^{-i}), \tilde{\theta})$ for all $k = 1, \dots, M^i$. 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 $k = 1, \dots, M^i$ 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 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 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

$$\begin{aligned}
U^i(s^i, s^i, s^{-i}) - U^i(\sigma^i, s^i, s^{-i}) &= \sum_{\tilde{s}^i \in \Gamma(\sigma^i)} \sigma^i(\tilde{s}^i) \left\{ \sum_{\theta \in \Theta} Q(\theta) [u^i((s^i, s^{-i}), \theta) - u^i((\tilde{s}^i, s^{-i}), \theta)] \right. \\
&+ \left. \sum_{\theta \in \Theta} Q(\theta) \sum_{\tilde{\theta} \in \Theta} Q(\tilde{\theta}) \left[\mu \left(u^i((s^i, s^{-i}), \theta) - u^i((s^i, s^{-i}), \tilde{\theta}) \right) - \mu \left(u^i((\tilde{s}^i, s^{-i}), \theta) - u^i((s^i, s^{-i}), \tilde{\theta}) \right) \right] \right\}.
\end{aligned} \tag{A.14}$$

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 $L^1(s_2^1, \sigma^2)$ the payoff lotteries for player 1 that are induced if he plays s_1^1 and s_2^1 , respectively. Since $\sigma^1(s_2^1) = 1 - \sigma^1(s_1^1)$, the utility of player one of playing σ^1 when expecting to play $\hat{\sigma}^1$ is given by:

$$\begin{aligned}
&U^1(\sigma^1, \hat{\sigma}^1, \sigma^2) \\
&= \sigma^1(s_1^1)E[L^1(s_1^1, \sigma^2)] + (1 - \sigma^1(s_1^1))E[L^1(s_2^1, \sigma^2)] + GL^1(\sigma^1, \hat{\sigma}^1, \sigma^2) \\
&= \sigma^1(s_1^1)E[L^1(s_1^1, \sigma^2)] + (1 - \sigma^1(s_1^1))E[L^1(s_2^1, \sigma^2)] \\
&+ \sigma^1(s_1^1)\hat{\sigma}^1(s_1^1)GL^1(s_1^1, s_1^1, \sigma^2) + (1 - \sigma^1(s_1^1))\hat{\sigma}^1(s_1^1)GL^1(s_2^1, s_1^1, \sigma^2) \\
&+ \sigma^1(s_1^1)(1 - \hat{\sigma}^1(s_1^1))GL^1(s_1^1, s_2^1, \sigma^2) + (1 - \sigma^1(s_1^1))(1 - \hat{\sigma}^1(s_1^1))GL^1(s_2^1, s_2^1, \sigma^2)
\end{aligned}$$

Taking the derivative with respect to $\sigma(s_1^1)$ yields:

$$\begin{aligned}
&\frac{\partial U(\sigma^1, \hat{\sigma}^1, \sigma^2)}{\partial \sigma^1(s_1^1)} \\
&= E[L^1(s_1^1, \sigma^2)] - E[L^1(s_2^1, \sigma^2)] + \hat{\sigma}^1(s_1^1)GL^1(s_1^1, s_1^1, \sigma^2) - \hat{\sigma}^1(s_1^1)GL^1(s_2^1, s_1^1, \sigma^2) \\
&+ (1 - \hat{\sigma}^1(s_1^1))GL^1(s_1^1, s_2^1, \sigma^2) - (1 - \hat{\sigma}^1(s_1^1))GL^1(s_2^1, s_2^1, \sigma^2) \\
&= E[L^1(s_1^1, \sigma^2)] - E[L^1(s_2^1, \sigma^2)] + GL^1(s_1^1, s_2^1, \sigma^2) - GL^1(s_2^1, s_2^1, \sigma^2) \\
&+ \hat{\sigma}^1(s_1^1)[GL^1(s_1^1, s_1^1, \sigma^2) - GL^1(s_2^1, s_1^1, \sigma^2) - GL^1(s_1^1, s_2^1, \sigma^2) + GL^1(s_2^1, s_2^1, \sigma^2)]
\end{aligned} \tag{A.15}$$

Suppose that s_1^1 and s_2^1 are not redundant for all $\sigma^2 \in \Sigma^2$. By Lemma 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

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

such that $\partial U(\sigma^1, \hat{\sigma}^1, \sigma^2)/\partial \sigma^1(s_1^1) \geq 0$ if and only if $\hat{\sigma}^1(s_1^1) \geq h(\sigma^2(s_1^2))$. 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 4.)

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 $\sigma_{\text{II}}^2 \in (0, 1)$ and $\sigma_{\text{I}}^2 \in [0, \sigma_{\text{II}}^2]$ such that $\sigma_{\text{II}}^2 = \min_{\sigma^2(s_1^2) \in (0, 1)} \{\sigma^2(s_1^2) | h(\sigma^2(s_1^2)) = 0\}$ and $\sigma_{\text{I}}^2 = \max_{\sigma^2(s_1^2) \in [0, \sigma_{\text{II}}^2]} \{\sigma^2(s_1^2) | h(\sigma^2(s_1^2)) = 1\}$. Since $h(\sigma^2(s_1^2))$ is a C^1 function, the set $\{(x, 0) | x \in [0, \sigma_{\text{II}}^2]\} \cup \{(x, h(x)) | x \in [\sigma_{\text{I}}^2, \sigma_{\text{II}}^2]\} \cup \{(x, 1) | x \in [\sigma_{\text{I}}^2, \sigma_{\text{II}}^2]\} \subseteq \mathcal{R}$ is connected. If $h(\sigma^2(s_1^2)) \leq 1$ for all $\sigma^2(s_1^2) \geq \sigma_{\text{II}}^2$, then the set $\mathcal{L} = \{(x, 0) | x \in [0, \sigma_{\text{II}}^2]\} \cup \{(x, h(x)) | x \in [\sigma_{\text{I}}^2, \sigma_{\text{II}}^2]\} \cup \{(x, 1) | x \in [\sigma_{\text{I}}^2, 1]\}$ 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 $\sigma^2(s_1^2) \in (\sigma_{\text{II}}^2, 1]$, then there exists $\sigma_{\text{IV}}^2 \in (\sigma_{\text{II}}^2, 1]$ and $\sigma_{\text{III}}^2 \in (\sigma_{\text{II}}^2, \sigma_{\text{IV}}^2)$ such that $\sigma_{\text{IV}}^2 = \min_{\sigma^2(s_1^2) \in (\sigma_{\text{II}}^2, 1]} \{\sigma^2(s_1^2) | h(\sigma^2(s_1^2)) = 1\}$ and $\sigma_{\text{III}}^2 = \max_{\sigma^2(s_1^2) \in [\sigma_{\text{II}}^2, \sigma_{\text{IV}}^2]} \{\sigma^2(s_1^2) | h(\sigma^2(s_1^2)) = 0\}$. The set $\{(x, 0) | x \in [0, \sigma_{\text{II}}^2]\} \cup \{(x, h(x)) | x \in [\sigma_{\text{I}}^2, \sigma_{\text{II}}^2]\} \cup \{(x, 1) | x \in [\sigma_{\text{I}}^2, \sigma_{\text{IV}}^2]\} \cup \{(x, h(x)) | x \in [\sigma_{\text{III}}^2, \sigma_{\text{IV}}^2]\} \cup \{(x, 0) | x \in [\sigma_{\text{III}}^2, \sigma_{\text{IV}}^2]\} \subseteq \mathcal{R}$ is a connected set. If $h(\sigma^2(s_1^2)) \geq 0$ for all $\sigma^2(s_1^2) \geq \sigma_{\text{IV}}^2$, the set $\mathcal{L} = \{(x, 0) | x \in [0, \sigma_{\text{II}}^2]\} \cup \{(x, h(x)) | x \in [\sigma_{\text{I}}^2, \sigma_{\text{II}}^2]\} \cup \{(x, 1) | x \in [\sigma_{\text{I}}^2, \sigma_{\text{IV}}^2]\} \cup \{(x, h(x)) | x \in [\sigma_{\text{III}}^2, \sigma_{\text{IV}}^2]\} \cup \{(x, 0) | x \in [\sigma_{\text{III}}^2, 1]\} \subseteq \mathcal{R}$ 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 $\sigma^2(s_1^2) \in (\sigma_{\text{IV}}^2, 1)$, we can proceed in the same way as we did from σ_{II}^2 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 $(1, R^1(1))$.

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_{II}^2 .

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




Figure 4: Illustration of the construction of the set \mathcal{L}

$h(\sigma^2(s_1^2)) \geq 1$ (≤ 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. \square

Proof of Proposition 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 $\chi((\hat{s}^{-i}, \hat{\theta}) | \sigma^{-i}) := Q(\hat{\theta}) (\prod_{j \neq i} \sigma^j(\hat{s}^j))$ 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

$$U^i(\tilde{s}^i, \tilde{s}^{-i}, \sigma^{-i}) = \sum_{(\hat{s}^{-i}, \hat{\theta}) \in \mathcal{X}} \chi((\hat{s}^{-i}, \hat{\theta}) | \sigma^{-i}) u^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta}) - \frac{\eta(\lambda - 1)}{2} \sum_{(\hat{s}^{-i}, \hat{\theta}) \in \mathcal{X}} \chi((\hat{s}^{-i}, \hat{\theta}) | \sigma^{-i}) \sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}} \chi((\tilde{s}^{-i}, \tilde{\theta}) | \sigma^{-i}) |u^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta}) - u^i((\tilde{s}^i, \tilde{s}^{-i}), \tilde{\theta})|. \quad (\text{A.16})$$

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.16) yields

$$\frac{dU^i(\tilde{s}^i, \tilde{s}^{-i}, \sigma^{-i})}{du^i((\tilde{s}^i, \hat{s}^{-i}), \hat{\theta})} = \chi((\hat{s}^{-i}, \hat{\theta}) | \sigma^{-i}) \left\{ 1 - \eta(\lambda - 1) \left[\sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}_+(\hat{s}^{-i}, \hat{\theta})} \chi((\tilde{s}^{-i}, \tilde{\theta}) | \sigma^{-i}) - \sum_{(\tilde{s}^{-i}, \tilde{\theta}) \in \mathcal{X}_-(\hat{s}^{-i}, \hat{\theta})} \chi((\tilde{s}^{-i}, \tilde{\theta}) | \sigma^{-i}) \right] \right\}. \quad (\text{A.17})$$

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.17) 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 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 $(\tilde{s}^i, \tilde{s}^{-i})$ different from (s^i, s^{-i}) . Since $(\tilde{s}^i, \tilde{s}^{-i})$ 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 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 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. \square

B. FOR ONLINE PUBLICATION

Proof of Proposition 6. We will show that the results from Proposition 1, Proposition 2, Proposition 4, Corollary 1 and Proposition 5 remain to hold in turn:

Regarding Proposition 1 for multidimensional outcomes:

(i) The coefficients $q_{km}^i(\sigma^{-i})$ for multidimensional outcomes differ from their counterparts for one-dimensional outcomes only in the sense that every comparison 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

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

holds, which is the analogue to (A.8) for multidimensional outcomes. Let $\Lambda_r^i(u) = \{(s, \theta) \in S \times \Theta \mid u_r^i(s, \theta) = u\}$ denote the set of (s, θ) combinations that result in some specific payoff $u_r \in \mathcal{U}_r^i$ 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 $P^i(u_r | \sigma) = \sum_{(s, \theta) \in \Lambda_r^i(u)} Q(\theta) \prod_{j=1}^I \sigma^j(s^j)$ and (B.1) can be rewritten equivalently as

$$\begin{aligned}
& \sum_{r=1}^R \left[\frac{1}{2} \sum_{u_r \in \mathcal{U}_r^i} \sum_{\tilde{u}_r \in \mathcal{U}_r^i} P^i(u_r | (\bar{\sigma}^i, \sigma^{-i})) \cdot P^i(\tilde{u}_r | (\bar{\sigma}^i, \sigma^{-i})) \cdot |u_r - \tilde{u}_r| \right. \\
& - \sum_{u_r \in \mathcal{U}_r^i} \sum_{\tilde{u}_r \in \mathcal{U}_r^i} P^i(u_r | (\bar{\sigma}^i, \sigma^{-i})) \cdot P^i(\tilde{u}_r | (\sigma_{\bar{x}}^i, \sigma^{-i})) \cdot |u_r - \tilde{u}_r| \\
& \left. + \frac{1}{2} \sum_{u_r \in \mathcal{U}_r^i} \sum_{\tilde{u}_r \in \mathcal{U}_r^i} P^i(u_r | (\sigma_{\bar{x}}^i, \sigma^{-i})) \cdot P^i(\tilde{u}_r | (\sigma_{\bar{x}}^i, \sigma^{-i})) \cdot |u_r - \tilde{u}_r| \right] = 0
\end{aligned}$$

By Lemma 1, this holds true if and only if the lotteries over material utility outcomes induced by $\bar{\sigma}^i$ and $\sigma_{\bar{x}}^i$ 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 $\sigma_{\tilde{x}}^i$ 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 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 $\sigma_{\tilde{x}}^i$. Thus, the lottery over multidimensional outcomes induced by the pure strategy being played with probability zero under $\sigma_{\tilde{x}}^i$ is a linear combination of the lotteries that are induced by the other pure strategies being played with positive probability, implying redundancy of $\sigma_{\tilde{x}}^i$.

(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 1(ii).

Regarding Proposition 2 for multidimensional outcomes:

The proof for multidimensional outcomes equals the proof of Proposition 2 up to (A.13), 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.13) for multidimensional outcomes is

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

By Lemma 1, this last inequality holds if and only if

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

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 $r = 1, \dots, R$ 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 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 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 $r = 1, \dots, R$. 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 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 $r = 1, \dots, R$, 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 4(ii).

Regarding Corollary 1 for multidimensional outcomes:

(i) The result follows directly from the fact that Proposition 2 continues to hold for multidimensional payoffs.

(ii) The result follows from Corollary 1(i) together with Proposition 7(i).

Regarding Proposition 5 for multidimensional outcomes:

According to the proof of Proposition 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 $r = 1, \dots, R$. 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 5 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 $r = 1, \dots, R$, which implies $U^i(s^i, \sigma^i, \sigma^{-i}) > U^i(\sigma^i, \sigma^i, \sigma^{-i})$ for all $\sigma^i \in \Sigma^i \setminus \{s^i\}$. □

Proof of Proposition 7. (i) The proof is identical to the corresponding proof of Proposition 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. □