INNOCENZO GASPARINI INSTITUTE FOR ECONOMIC RESEARCH

Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Path Agreements and Off-Path Uncertainty

Pierpaolo Battigalli and Emiliano Catonini

Working Paper n. 736

This Version: November 26, 2026

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
http://www.igier.unibocconi.it

The opinions expressed in the working papers are those of the authors alone, and not those of the Institute, which takes non institutional policy position, nor those of CEPR, NBER or Università Bocconi.

Path Agreements and Off-Path Uncertainty*

Pierpaolo Battigalli† Emiliano Catonini‡

This draft: November 26, 2026

Abstract

We consider a multistage game where players have reached a pre-play, non-binding agreement on which path to follow. We study the credibility of the agreement, that is, incentive compatibility under forward-induction reasoning, when players are averse to the residual uncertainty that a deviation entails. We show that the set of credible paths expands as uncertainty aversion (ambiguity aversion, or risk aversion under ambiguity neutrality) increases. This is true also when only an upper bound on risk and ambiguity aversion is commonly known. While credibility does not imply necessary compliance with the agreed-upon path, we show that, in finite games, when players perceive sufficient uncertainty and are sufficiently risk or ambiguity averse, compliance with the path is guaranteed. This supports the view that aversion to the uncertainty that deviations entail can be the fundamental reason for staying on the agreed-upon path, rather than precise off-path beliefs imposed by Nash equilibrium and its standard refinements.

KEYWORDS: path agreements, comparative uncertainty aversion, credibility, self-enforceability, forward induction, monotonicity.

1 Introduction

A classical motivation for equilibrium analysis is that players can communicate before the game starts and reach a non-binding agreement about how to play. Nash equilibrium and refinements model this scenario with a profile of strategies that is supposedly and informally deemed self-enforcing. In dynamic games, there are at least two reasons why this may not be an accurate

*We thank Andres Perea, Federico Golonia, Laura Lo Schiavo, Fabio Macceromni, Silvia Meneghesso, William Miller, and the participants to the 5th UNIBG Industrial Organization Winter Symposium, and brown bag seminars at Singapore Management University, Higher School of Economics (ICEF), and NYU Shanghai for their comments. This project is funded by the European Union (ERC, TRAITS-GAMES, 101142844). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

†Bocconi University and IGIER, pierpaolo.battigalli@unibocconi.it

‡New York University in Shanghai, Department of Economics, emiliano.catonini@gmail.com

representation of the strategic situation that players are facing. First, real-life agreements are often incomplete: Players may fail to reach an agreement about what to do once someone defects, or may not even try to do it, because they anticipate that they will not trust each other anymore after a deviation, or because even discussing the possibility of a deviation is taboo. Second, in a game without chance moves, if a non-binding agreement is represented by a pure equilibrium strategy profile, players' attitudes toward uncertainty do not matter. If instead the agreement does not specify off-path behavior, players are likely to perceive strategic uncertainty about the behavior of co-players after a deviation. In this case, their attitudes toward uncertainty matter and their aversion to uncertainty may affect the incentive to deviate.[1]

Here we take the view that strategic uncertainty regarding off-path behavior and aversion towards it play a key role for the decision to comply with the agreement. Hence, we study the behavior of players who only agree on a path to follow, and we do not introduce belief restrictions about off-path behavior beyond those implied by strategic reasoning based on forward induction. Indeed, the latter may reduce the perceived uncertainty and challenge the incentive to comply. Specifically, forward-induction considerations may lead to the conclusion that a deviation will certainly be profitable. If this is *not* the case, as in Catonini (2021), we will say that the path (agreement) is "**credible**."

Focusing on an environment without objective risk, we can interpret uncertainty aversion so as to encompass both risk and ambiguity aversion.[2] We represent an increase in uncertainty aversion by means of concave and strictly increasing transformations of players' payoff functions. With this, we show that *the set of credible paths expands as uncertainty aversion increases.* First, we derive this result under common knowledge of uncertainty attitudes. At the other extreme, in a kind of "belief-free" perspective, we show that the set of credible paths expands also when only an upper bound on players' uncertainty aversions is commonly known, and this upper bound increases. Finally, we show that, in a finite game, if players' beliefs are restricted to capture a minimal level of perceived uncertainty, an increase in uncertainty aversion makes a credible path "self-enforcing," in the sense that players have no incentive to deviate for *any* refined belief. These results support the view that players' aversion towards the off-path uncertainty is a valid explanation for compliance.

The idea that the uncertainty entailed by deviations can be crucial to explain compliance with an expected behavior has already been investigated in the context of self-confirming equilibrium analysis (Fudenberg & Levine 1993, Battigalli 1987).[3] Battigalli *et al.* (2015) define self-confirming equilibrium with ambiguity aversion for a simultaneous-move game played recurrently with random

[1] We do not endorse the view that an agreement to play a subgame perfect equilibrium would be credible and eliminate strategic uncertainty, because such an agreement may unravel when deviations are interpreted according to forward-induction reasoning as intentional attempts to increase payoffs.

[2] Our analysis relies on dynamic consistency. Ambiguity aversion – the reluctance to make choices involving unknown risks – may yield relevant dynamic inconsistencies. Battigalli *et al.* (2019) show when this is not the case.

[3] See also related observations in Greenberg (2000) and Riedel & Sass (2014).

matching of agents from different populations. They show that the set of self-confirming equilibria expands as ambiguity aversion increases. Battigalli *et al.* (2019) extend the analysis to sequential games, and show that, because of dynamic inconsistencies of preferences, the monotonicity result of Battigalli *et al.* (2015) does not hold through. However, in terms of outcomes, monotonicity can be re-established under restrictive conditions on the game and on the nature of the equilibrium.

In our context, the monetary payoff functions can be plausibly assumed to be commonly known, and players can use this fact to reason strategically about the behavior of co-players. Thus, credible path agreements are – technically – a kind of rationalizable self-confirming equilibrium,[4] but we are not necessarily trying to explain stability of behavior in a situation of recurrent strategic interaction, as our analysis applies (also) to games played just once under a pre-play, non-binding agreement. Greenberg *et al.* (2008) put forward a somewhat similar notion of "mutually acceptable course of action." Unlike them, we follow Catonini (2021) in assuming that beliefs after deviations from the agreed-upon behavior are shaped by forward-induction reasoning.

In Section 4, we argue that our analysis encompasses both risk and ambiguity aversion. As in Battigalli *et al.* (2019), we consider players with smooth ambiguity preferences (Klibanoff *et al.* 2005) over actions at each information set, given their beliefs about others and continuation plans. Battigalli *et al.* (2019) observed that, in the absence of chance moves, when every agent in each population follows the same (pure) strategy, ambiguity and risk attitudes are indistinguishable and, as a consequence, players' preferences are dynamically consistent. For us, each player-role represents not a population, but an individual agent; therefore, we can treat our players as if they were subjective expected utility maximizers with a von Neumann-Morgenstern utility function that summarizes the attitudes towards risk and ambiguity. We model players' beliefs with conditional probability systems and restrict them to initially assign probability 1 to the strategies of the co-players that comply with the agreed-upon path. Let Δ denote the profile of subsets of beliefs (one subset for each player) that satisfy this restriction. Call a strategy "**justifiable**" if it is optimal (in a sequential sense) under some belief in the reference belief set. With this, we can compute the justifiable strategies of players who tentatively believe in the path, i.e., given that their beliefs belong to Δ. We first consider the simple case in which risk and ambiguity attitudes are commonly known. Then, players can anticipate what strategies of the co-players are justifiable under the belief in the path, and refine their own beliefs with strategic reasoning, ascribing to the co-players the highest degree of "strategic sophistication" consistent with what they observe. The iterated removal of unjustifiable strategies is called Strong-Δ-Rationalizability.[5] So, Strong-Δ-Rationalizability captures the behavior of rational players who believe in the agreed-upon path, *believe as long as possible* that the co-players are rational and believe in the path, and so on. See

[4] Assuming that actions are observed also ex post, when the game is over.

[5] See Battigalli (2003) and Battigalli & Siniscalchi (2003). More recently, the term "directed rationalizability" has been used to refer to the map from the given belief restrictions to the resulting solution. See, e.g., Battigalli & Catonini (2026), and the textbook by Battigalli, Catonini & De Vito (2025), chapters 8 and 15.

Battigalli & Siniscalchi (2007), Battigalli & Prestipino (2013) and Battigalli & Tebaldi (2019) for a justification based on a rigorous formal analysis of interactive conditional beliefs. We say that a path is **credible** if Strong-Δ-Rationalizability yields a non-empty set, which means that, for every player, staying on path is compatible with strategic reasoning.[6]

In our framework, the analysis of Battigalli *et al.* (2015) applies, and it implies that if a player who believes in the path has no incentive to deviate for some given belief about how the co-players would react, this will be the case also after an increase of her uncertainty aversion. A dynamic extension of the results of Battigalli *et al.* (2016) and Weinstein (2016) further guarantees that the whole set of sequentially justifiable strategies can only expand after an increase in uncertainty aversion. This means that also the possible reactions to deviations can only expand, preserving players' incentives to stay on path after one step of reasoning. At this point, one may hope that a simple inductive argument would yield the desired monotonicity result, but this is not the case. The conceptual reason is that "believing as long as possible" that a proposition is true does *not* imply "believing as long as possible" that a *weaker* (more permissive) proposition is true, an issue that we formally address in Section 2. Specifically, consider a deviation that, at some step of reasoning, is not profitable for any belief under the baseline attitudes. Under increased uncertainty aversion, for some belief over the larger set of possible reactions, the deviation could still be perceived as profitable. Then, after this deviation, players can keep refining their beliefs with forward-induction reasoning. We will show by example that this may even lead to a completely different set of rationalizable reactions to the deviation. Yet, perhaps surprisingly, *there will always remain reactions that deter the deviation.* Proving this fact is the hardest challenge for this paper. The trick is to view Strong-Δ-Rationalizability under higher uncertainty aversion as a slow, *unfinished* order of elimination under lower uncertainty aversion, where one "forgets" to eliminate the strategies that can be justified under higher but not under lower aversion. Then, we can borrow intuitions and techniques from the papers that study strategic reasoning under non-monotone reduction operators. In particular, we will extend to games with infinite horizon and to the case of unknown uncertainty attitudes the techniques of Battigalli & Catonini (2026).[7] In a nutshell, the intuition behind these techniques is the following: Slowing down the elimination procedure (i.e., increasing uncertainty aversion) cannot select reactions to any given deviation that *all* make the deviation profitable, for if it did, a subset of those reactions would have also survived the baseline elimination procedure (i.e., under the baseline attitudes).

The relationship between uncertainty attitudes and the ability to coordinate with non-binding and possibly tacit agreements has been studied theoretically and experimentally under different

[6] In Catonini (2021), credibility of an agreement is defined with respect to another elimination procedure, called Selective Rationalizability (Catonini 2019). Catonini (2020) shows that the two procedures are equivalent for path agreements. See Section 2 for details.

[7] Which in turn are related to earlier work by Chen & Micali (2013), Perea (2018a, 2018b) and Catonini (2020). See also Perea (2025).

assumptions. Bernhardt & Rastad (2016) study a repeated oligopoly with stochastic demand, fixed costs, and risk-averse firms. They find that greater risk aversion makes collusion harder to sustain in recessions but easier in booms. Unlike us, by assumption they consider on-path risk. The mechanism works through the curvature of utility over stochastic profits and the relative attractiveness of deviations, rather than through uncertainty induced by leaving an agreed-upon path. Rojas & Vergara (2021) introduce neo-additive ambiguous beliefs into infinitely repeated social dilemmas. Players distrust their predictions of opponents' behavior, and pessimism represents ambiguity aversion. Since ambiguity concerns whether the opponent will adhere to the cooperative strategy, aversion to ambiguity makes the cooperative path less attractive, and this may raise the critical discount factor for cooperation. We instead assume that players trust path predictions. In experimental work, Sabater-Grande & Georgantzis (2002) elicit subjects' risk attitudes and then have matched subjects play repeated prisoner's dilemmas. More risk-averse subjects cooperate less, and the negative relationship becomes stronger when continuation is stochastic rather than deterministic.[8] Their interpretation is that cooperation is risky: it exposes a player to exploitation by the other player. This is reasonable given that subjects in the experiment cannot reach a pre-play agreement, as we instead assume.

Many other papers have tackled the issue of coordination under forward-induction reasoning or under uncertainty aversion, but, to the best of our knowledge, sticking to (sequential) equilibrium. With complete information about payoffs, standard equilibrium analysis eliminates the perception of ambiguity, because players are assumed to hold Dirac beliefs about the (possibly random) behavior of others both on and off path.[9] Instead we model strategic reasoning allowing for genuine uncertainty about off-path behavior even under complete information. Furthermore, we consider an environment that allows us to consider risk and ambiguity attitudes together. A vast literature, stemming from the seminal contribution of Kohlberg and Mertens (1986), has refined sequential equilibrium by means of forward-induction considerations. Cho & Kreps (1987), Banks & Sobel (1987), and Govindan & Wilson (2009) motivate their solution concepts with the idea that deviations are interpreted as attempts to improve upon the equilibrium payoff.[10] This means that the deviator is assumed to believe in the equilibrium path, but not in the equilibrium threats. But this destroys the motivation for the off-path coordination of beliefs assumed by the sequential equilibrium concept; therefore, players are exposed to strategic uncertainty. Our analysis can also be interpreted as a forward-induction test for the threats that sustain an equilibrium path,[11] and the threats that pass the test need not be consistent with sequential equilibrium.

[8] In a repeated oligopoly experiment, Anderson *et al.* (2012) reach a similar empirical conclusion.

[9] Incomplete information allows for the perception of ambiguity concerning the types of others and the predictive probabilities of others' actions.

[10] These intuitions have indeed been confirmed through the use of Strong-Δ-Rationalizability with path restrictions by Battigalli & Siniscalchi (2003) for the (Iterated) Intuitive Criterion, and by Battigalli & Catonini (2024) for Divine Equilibrium (with additional belief restrictions related to independence).

[11] In case of ambiguity aversion, only "ambiguous threats" may pass the test.

Another stream of literature has introduced ambiguity aversion in equilibrium analysis. Lo (1999) analyzes a notion of Nash equilibrium whereby players are mini-maximizers given belief sets in the sense of Gilboa & Schmeidler (1989). Riedel & Sass (2014) apply the maxmin criterion to analyze equilibria of simultaneous-move games where players may delegate the selection of actions to random devices with unknown outcome probabilities ("Ellsberg urns"). Hanany *et al.* (2020) introduce a notion of perfect Bayesian equilibrium with the smooth ambiguity criterion. Eichberger *et al.* (2018) analyze a notion of equilibrium with ambiguity aversion in sequential games that combines Choquet-expected utility maximization (Schmeidler 1989) with a generalized form of Bayesian updating.[12]

For the analysis of a bargaining game, Friedenberg (2019) introduces the notion of "path strategic certainty", that is, the (correct) belief that the co-player will not deviate from the given path. Like us, Friedenberg introduces this element of partial coordination in the context of common strong belief in rationality (Battigalli & Siniscalchi 2002), rather than assuming complete coordination on a profile of equilibrium strategies. Differently from us, Friedenberg's kind of forward-induction reasoning is not based on the belief in the given path.

The rest of the paper is organized as follows. Section 2 introduces our theoretical framework, illustrating it with a running example. Section 3 contains the main results.

2 Framework

Description of the game form We consider a multistage game form with *observed actions* chosen from a *finite* set, *no chance moves*, either finite or infinite horizon, and monetary payoffs.[13] Our analysis also covers the infinite repetition with perfect monitoring of a finite one-period game. We interpret the absence of chance moves also as an assumption that players cannot delegate their choices to randomization devices. Finiteness of the action sets guarantees compact metric strategy sets in the natural "discounting metric" even if the horizon is infinite. Unless we specify otherwise, for technical details related to the analysis of multistage games with infinite horizon we refer to Battigalli (2003).

There is a finite set of players I. Each player i has a *finite* set of actions A_i that she can choose at some history of the game. At every non-terminal history, all players simultaneously choose one of their available actions. Thus, a history is a sequence of actions profiles. Let H denote the set of **non-terminal histories**, i.e., the finite sequences of action profiles allowed by the rules of the game that do not terminate the game, including the initial history (empty history) $h = \varnothing$. Note that

[12] See also, Eichberger and Kelsey (2011) and Eichberger *et al.* (2009).

[13] We call such structures "game forms" because they only describe the rules of the game, not players' attitudes toward uncertainty. Our framework is broadly consistent with the textbook by Osborne & Rubinstein (1994). See also the textbook by Battigalli, Catonini & De Vito (2025). The assumption of observed actions is made for simplicity. We can show that the analysis can be extended to sequential games played by agents with perfect recall.

H is *countable*, i.e., finite or – in the infinite-horizon case – denumerable. For each non-terminal history $h \in H$ and player $i \in I$, $A_i(h)$ denotes the finite set of **actions available** to i **at** h. If player i is *not active* at h, $A_i(h)$ is a *singleton*. We let Z denote the set of **terminal histories**, or **paths**, of the game (including infinite histories if there is an infinite horizon) and we endow the set of all histories $\bar{H} = H \cup Z$ with the weak (respectively, strict) "prefix-of" relation $\preceq$ ($\prec$), which makes $\bar{H}$ a *tree* with *root* $\varnothing$. We let $\pi_i : Z \to \mathbb{R}$ denote the **monetary payoff function** of player i, with the interpretation that in multi-period games with end-of-period payoffs, $\pi_i(z)$ represents the discounted present value of the stream of payoffs on path z. This justifies the assumption that each π_i is *continuous* and *bounded* (given the standard "discounting" metric on the set of infinite paths of play).[14] Players only care about their own monetary payoff, but they may be risk averse; thus, $\pi_i(z)$ does not represent the von Neumann-Morgenstern utility assigned by i to z.[15]

A **strategy** of player i is a map $s_i \in S_i = \times_{h \in H} A_i(h)$ that assigns an action $a_i \in A_i(h)$ to every non-terminal history $h \in H$. We let $S = \times_{j \in I} S_j$ denote the set of strategy profiles, and we let $S_{-i} = \times_{j \in I \setminus \{i\}} S_j$ denote the set of strategy profiles of the co-players of i. The **path function** $\zeta : S \to Z$ associates each strategy profile with the induced terminal history. In infinite-horizon games the path function ζ is continuous. For each $h \in \bar{H}$, we let

$$S_i(h) = \{s_i \in S_i : \exists s_{-i} \in S_{-i}, h \preceq \zeta(s_i, s_{-i})\}$$

denote the set of strategies of player i that are consistent with history h. The set of strategy profiles inducing any history h is the cross-product

$$\{s \in S : h \preceq \zeta(s)\} = S(h) = \times_{i \in I} S_i(h);$$

thus, for each $i \in I$, the set of co-players' strategy profiles consistent with h is $S_{-i}(h) = \times_{j \in I \setminus \{i\}} S_j(h)$, and $S(h) = S_i(h) \times S_{-i}(h)$. All these sets of strategies, or strategy profiles are compact and, for each finite h, clopen.[16] Conversely, we let $H(s_i) = \{h \in H : s_i \in S_i(h)\}$ denote the set of non-terminal histories that are consistent with strategy s_i. Differently from the perfect and sequential equilibrium concepts, our analysis rests on the assumption that observed actions by co-players are always interpreted as *intentional*, even when they deviate from the agreed-upon path. In other words, it is transparent[17] players' plans coincide with their history-dependent behavior. Thus, the unresolved uncertainty about co-players' behavior at history h is represented by set $S_{-i}(h)$.

[14] In multi-period infinite-horizon games, this follows from common bounds on the per-period monetary payoff functions and standard exponential discounting.

[15] Therefore, the primitive elements listed so far describe a **game form** with monetary payoffs rather than a game in the technical sense. Players' preferences over monetary lotteries will be introduced to describe their attitudes toward risk and uncertainty.

[16] That is, simultaneously closed and open, as they represent observations about earlier choices in finite sets.

[17] True and commonly believed.

Example 1 *We illustrate our analysis through the following game form with monetary payoffs.*

A\B	W	E
N	9, 9	·–
S	0, 0	9, 9

$\longrightarrow$

A\B	L	C	M	R
T	4, 2	4, 5	4, 16	4, 0
U	4, 2	4, 5	4, 0	4, 16
D	0, 2	16, 1	0, 16	0, 0
B	16, 2	0, 1	0, 0	0, 16

Figure 1

At the initial history, Ann and Bob choose between two actions: $A_a(\varnothing) = \{\mathrm{N}, \mathrm{S}\}$, $A_b(\varnothing) = \{\mathrm{W}, \mathrm{E}\}$. *If they play* (N, W), (S, E), *or* (S, W), *the game ends, thus* $((\mathrm{N}, \mathrm{W})), ((\mathrm{S}, \mathrm{W})), ((\mathrm{S}, \mathrm{E})) \in Z$. *If they play* (N, E), *the game moves to the second stage, hence* $H = \{\varnothing, ((\mathrm{N}, \mathrm{E}))\}$. *At history* $((\mathrm{N}, \mathrm{E}))$, *Ann and Bob choose between 4 actions, and then the game ends with one of the 16 terminal histories of the form* $((\mathrm{N}, \mathrm{E}), (a, b)) \in Z$ *(*$a \in \{\mathrm{T}, \mathrm{U}, \mathrm{B}, \mathrm{D}\}$, $b \in \{\mathrm{L}, \mathrm{C}, \mathrm{M}, \mathrm{R}\}$*). The monetary payoffs are specified in Figure 1.*

Attitudes toward uncertainty Let $B \subseteq \mathbb{R}$ denote a bounded interval containing the range $\pi_i(Z)$ of monetary payoffs. Player i's attitudes towards uncertainty are represented by a continuous and strictly increasing utility function $v_i : B \to \mathbb{R}$, which we assume to be *concave*. An increase in the concavity of v_i (by means of a concave and strictly increasing transformation) represents an increase in uncertainty aversion. To ease reading, one may interpret v_i as a von Neumann-Morgenstern utility function representing attitudes toward risk, although we consider an environment with mere strategic uncertainty and without objective risk. In Section 4 we show that our analysis applies also to ambiguity aversion. For each player i, we obtain the (continuous) **payoff function** $u_i = v_i \circ \pi_i : Z \to \mathbb{R}$, that is, the utility of paths of play.

Beliefs, sequential optimality, and justifiability We assume that each player has a conditional belief over the co-players' strategies at every non-terminal history, and that this belief is derived by updating the belief at the preceding history whenever possible. That is, we assume that each player i has an array of conditional beliefs $\mu^i = (\mu^i(\cdot|h))_{h \in H} \in \times_{h \in H} \Delta(S_{-i})$ where $\mu^i(S_{-i}(h)|h) = 1$ for each $h \in H$, and

$$\mu^i(E_{-i}|h') \cdot \mu^i(S_{-i}(h')|h) = \mu^i(E_{-i}|h)$$

for all $h' \succ h$ and $E_{-i} \subseteq S_{-i}(h')$.[18] Battigalli *et al.* (2023) call such belief systems "forward consistent" and show that they are equivalent to standard conditional probability systems (CPSs)

[18] In infinite-horizon games, $\Delta(S_{-i})$ is the set of Borel probability measures on the compact metric space S_{-i}, endowed with the topology of weak convergence of measures.

à la Renyi (1955) for the analysis of weak sequential optimality and strong Δ-rationalizability, which we introduce below. Thus, for brevity, we simply refer to them as "CPSs". Let $\Delta^H(S_{-i})$ denote the set of all CPSs of i.[19]

For any (measurable) $E_{-i} \subseteq S_{-i}$, let

$$H(E_{-i}) = \{h \in H : \exists s_{-i} \in E_{-i}, s_{-i} \in S_{-i}(h)\}$$

denote the set of nonterminal histories consistent with co-players playing strategies in E_{-i}. Following Battigalli & Siniscalchi (2002), we say that a CPS μ^i **strongly believes** a strategic event $E_{-i} \subseteq S_{-i}$ if $\mu^i(E_{-i}|h) = 1$ for every $h \in H(E_{-i})$. With this, we let $\Delta^H_{\text{sb}}(E_{-i})$ denote the **set of all CPSs that strongly believe** E_{-i}. Whenever E_{-i} is compact, $\Delta^H_{\text{sb}}(E_{-i})$ is compact as well (see Battigalli & Tebaldi 2019).

It is important to observe that belief set $\Delta^H_{\text{sb}}(E_{-i})$ is *not monotone* in its argument E_{-i}. Specifically, let $E'_{-i} \subset E''_{-i}$(strict inclusion); then $H\left(E'_{-i}\right) \subseteq H\left(E''_{-i}\right)$ (weak inclusion) and $\mu^i \in \Delta^H_{\text{sb}}(E'_{-i})$ implies $\mu^i\left(E'_{-i}|h\right) = \mu^i\left(E''_{-i}|h\right) = 1$ for all $h \in H\left(E'_{-i}\right)$; if $H\left(E'_{-i}\right) = H\left(E''_{-i}\right)$ then $\Delta^H_{\text{sb}}(E'_{-i}) \subset \Delta^H_{\text{sb}}(E''_{-i})$; if instead $H\left(E'_{-i}\right) \subset H\left(E''_{-i}\right)$, $\Delta^H_{\text{sb}}(E'_{-i})$ and $\Delta^H_{\text{sb}}(E''_{-i})$ are not necessarily nested, because for any $h \in H\left(E''_{-i}\right) \setminus H\left(E'_{-i}\right)$ we may have $\mu^i\left(E''_{-i}|h\right) < 1$ implying $\mu^i \notin \Delta^H_{\text{sb}}(E''_{-i})$. This lack of monotonicity is what makes the proof of the main theorem hard.

Strategy s^*_i is **weakly sequentially optimal under** CPS μ^i if, for each $h \in H(s^*_i)$,

$$s^*_i \in \arg\max_{s_i \in S_i(h)} \mathbb{E}_{\mu^i(\cdot|h)}\left[u_i(\zeta(s_i, s_{-i}))\right].$$

Let $r_i(\mu^i)$ denote the set of strategies that are weakly sequentially optimal under μ^i. Compactness of S_i and continuity of path function ζ and of payoff (utility) function u_i guarantee that the weak sequential best reply correspondence $\mu^i \mapsto r_i(\mu^i)$ is *upper-hemicontinuous* and *nonempty compact-valued*.

The adverb "weakly" refers to requiring continuation optimality only for nonterminal histories consistent with the given strategy, which allows to apply such notion of optimality to *reduced* strategies.[20] For brevity, we henceforth say that each $s_i \in r_i(\mu^i)$ is "**optimal**" under μ^i. When we want to emphasize the uncertainty attitudes built in $u_i = v_i \circ \pi_i$, we say that s_i is "v_i-**optimal**" under μ^i.

Fix some set of CPSs $\Delta_i \subseteq \Delta^H(S_{-i})$. Strategy s^*_i is **justifiable** given Δ_i if it is optimal under some $\mu^i \in \Delta_i$.

[19] In finite games, $\Delta^H(S_{-i})$ is a compact subset of a Euclidean space. In the infinite-horizon games considered here, $\Delta^H(S_{-i})$ is a closed – hence, compact – subset of a compact metrizable space (see Battigalli & Siniscalchi 1999, Lemma 1).

[20] A **reduced strategy** of i is an element of the quotient set $S_i|\sim_i$, where $\sim_i$ is the realization equivalence relation on S_i. See the discussion in Battigalli *et al.* (2023).

Path agreements and strategic reasoning We first explain *informally* our assumptions about strategic reasoning by rational players, focusing for simplicity on the case of two players.[21] Next we introduce the solution concept that characterizes the behavioral implications of such assumptions, an iterated removal of unjustifiable strategies akin to the strong version of rationalizability for sequential games first put forward by Pearce (1984) and further analyzed and extended by Battigalli (1997, 2003) and Battigalli & Siniscalchi (2002, 2003). Battigalli (2003) coined the term "Strong Rationalizability" for this solution concept, because it relies on strong belief in rationality and it refines an alternative solution concept put forward by Ben Porath (1997).[22]

We consider rational players who have agreed to follow a path z. Thus, they initially believe that the co-player will not deviate from this path. Once a unilateral deviation by co-player j is observed, we assume that player i, if possible, still believes that player j is rational, initially believed that i would have complied with the agreed-upon path z, and reasons strategically. Thus, player i tries to rationalize j's deviation and infer her continuation plan by forward induction. In particular, player i interprets j's deviation as an attempt to improve her payoff with respect to $v_j(z)$. In this section, we assume that players' preferences are commonly known. Then, player i can determine the continuation plans of j that could potentially make the deviation profitable, and best reply to a belief that accounts for this. In turn, player j can (partially) anticipate i's reaction to the deviation. This rests on the hypothesis that player j strongly believes that player i is rational, initially believes that j would comply with the agreed-upon path z, and strongly believes that the same is true for j. In sum, we assume that players are rational and initially believe in the co-player's compliance, and that there is "common strong belief" of this.

The behavioral implications of these assumptions about rationality and strategic reasoning are characterized by **Strong Directed Rationalizability** (Battigalli 2003, Battigalli & Siniscalchi 2003), a solution concept that captures forward-induction reasoning under the assumption that some restrictions on players' beliefs hold and are commonly believed.[23] The assumed belief restrictions "direct" the solution toward a result (set of strategy profiles). Absent such restrictions, one obtains the aforementioned Strong Rationalizability concept. Here we assume each player believes that co-players' behavior complies with an agreed-upon path of play. Formally, fix a path $z \in Z$, and for each player i, let Δ_i^z denote the compact set of conditional probability systems μ^i that initially believe in $S_{-i}(z)$; that is, those that satisfy $\mu^i(S_{-i}(z)|\varnothing) = 1$. Thus, we posit belief restrictions

[21] Formally, such assumptions concern the relationship between behavior and hierarchies of conditional beliefs analyzed by the epistemic theory of sequential games; see Catonini (2019), Battigalli & Tebaldi (2019) and the relevant references therein.

[22] Battigalli (1997) and many successive articles call this solution concept "extensive form rationalizability." But such terminology has become ambiguous and inappropriate, because there are other versions of the rationalizability idea for sequential games represented in extensive form that are justified by rigorous epistemic foundations; see, e.g., Ben Porath (1997). The main result of Chen & Micali (2013) implies that, in finite games, strong rationalizability also refines "backward" versions of rationalizability (Perea 2014), although only in outcomes, not strategies.

[23] See, e.g., Battigalli & Prestipino (2013), Battigalli & Tebaldi (2019), and Catonini (2019).

represented by the profile $\Delta^z = (\Delta_i^z)_{i\in I}$ of subsets of CPSs. Given Δ^z, we define Strong-Δ^z-Rationalizability, abbreviated as "z-**Rationalizability**". For any (measurable) $E_{-i} \subseteq S_{-i}$, let $\Delta_{i,\text{sb}}^z(E_{-i}) = \Delta_i^z \cap \Delta_{\text{sb}}^H(E_{-i})$ denote the set of CPSs that initially believe in co-players' compliance with z and strongly believe E_{-i}. We define z-rationalizability recursively with the following iterated deletion of unjustifiable strategies:

- $(n = 0)$ for each $i \in I$, let $S_{i,z}^0 = S_i$;
- $(n \in \mathbb{N})$ for each $i \in I$ and $s_i \in S_{i,z}^{n-1}$, let $s_i \in S_{i,z}^n$ if and only if $s_i \in r_i(\mu^i)$ for some $\mu^i \in \Delta_{i,\text{sb}}^z(S_{-i,z}^{n-1})$;
- $(n = \infty)$ for each $i \in I$, let $S_{i,z}^\infty = \cap_{n>0} S_{i,z}^n$.

We call $S_{i,z}^\infty$ the set of z-**rationalizable** strategies of player i. Note that each $S_{i,z}^n$ is compact.

Path z may be such that, for some i and n, $\Delta_{i,\text{sb}}^z(S_{-i,z}^{n-1}) = \Delta_i^z \cap \Delta_{\text{sb}}^H(S_{-i,z}^{n-1}) = \emptyset$. In this case, $S_{i,z}^n = \emptyset$ as well. We say that path z is **credible** when $S_z^\infty \neq \emptyset$. Given our belief restrictions, the set of z-rationalizable strategy profiles is nonempty if and only if at least one profile of strategies that comply with the path survives every step of the elimination procedure. This means that staying on path is compatible with strategic reasoning for each player: even after refining beliefs about off-path behavior by forward induction, no unilateral deviation from the path will surely be profitable.

Our definition of credibility of a path relies on a simplified iterated deletion procedure. One simplification is that the CPSs justifying strategies in $S_{i,z}^n$ are only required to strongly believe $S_{-i,z}^{n-1}$, but not subsets $S_{-i,z}^k \supseteq S_{-i,z}^{n-1}$ obtained in previous steps $k < n-1$. Since strong belief is not monotone, strong belief in $S_{-i,z}^{n-1}$ does not by itself imply imply strong belief in $S_{-i,z}^k \supseteq S_{-i,z}^{n-1}$. Yet, the elimination procedure actually captures a "**best rationalization principle**" implying this.[24] Indeed, we can interpret each $S_{i,z}^k$ as the set of strategies of i that are consistent with "level k" of rationality and strategic sophistication given the agreed-upon path z. With this, for every $n \in \mathbb{N}\cup\{\infty\}$, $S_{i,z}^n$ is the set of strategies of i justified by beliefs in Δ_i^z ascribing to the co-players the highest level of rationality and strategic sophistication below n consistent with observed behavior.

Remark 1 *Fix a path $z \in Z$. For each player $i \in I$ and $n \in \mathbb{N}$,*

$$\begin{aligned} S_{i,z}^n &= \left\{s_i \in S_i : \exists \mu^i \in \Delta_i^z, \forall k < n, \forall h \in H\left(S_{-i,z}^k\right), \mu^i\left(S_{-i,z}^k | h\right) = 1, s_i \in r_i\left(\mu^i\right)\right\}, \\ S_{i,z}^\infty &= \left\{s_i \in S_i : \exists \mu^i \in \Delta_i^z, \forall k \in \mathbb{N}\cup\{\infty\}, \forall h \in H\left(S_{-i,z}^k\right), \mu^i\left(S_{-i,z}^k | h\right) = 1, s_i \in r_i\left(\mu^i\right)\right\}. \end{aligned}$$

In words, $s_i \in S_{i,z}^n$ (with $n \in \mathbb{N}\cup\{\infty\}$) if and only if $s_i \in r_i\left(\mu^i\right)$ for some $\mu^i \in \Delta_i^z$ that strongly believes $S_{-i,z}^k$ *for each* $k < n$, and also for $k = \infty$ if $n = \infty$. We will use a related result repeatedly.

[24] See, e.g., Battigalli & Siniscalchi 2002, Battigalli *et al.* 2025, and the relevant references therein.

Its proof is a rather straightforward adaptation to infinite games of similar results in the literature.[25]

Consider finite games. Catonini (2021) provides an analogous definition of credibility for a pre-play, non-binding agreement (such as the path-agreement considered here) based on another rationalizability procedure with belief restrictions: Selective Rationalizability. Differently from Strong Directed Rationalizability, Selective Rationalizability requires to keep the highest order of belief in rationality and strategic sophistication that is compatible with observed behavior, even when this requires to drop an order of belief in the belief restrictions.[26] The epistemic priority assigned to beliefs in rationality and strategic sophistication of all orders implies that selective rationalizability refines Strong Rationalizability (Pearce 1984, Battigalli 1997), while Strong Directed Rationalizability does not.[27] However, Catonini (2020) shows that the two solution concepts are outcome-equivalent when the first-order belief restrictions correspond to the belief in a specific path z.[28] So, we stick to z-Rationalizability, which is simpler from an algorithmic viewpoint. The following example illustrates the iterated deletion procedure.

Example 2 *Consider the game form with monetary payoffs of Example 1. Suppose players are risk and ambiguity neutral, that is, $u_i = \pi_i$ $(i = a, b)$. Suppose also that they agree on path $z = ((\mathrm{S}, \mathrm{E}))$. Then, we have*

$$\begin{aligned} \Delta_a^z &= \left\{\mu^a \in \Delta^H(S_b) : \mu^a(\mathrm{W}|\varnothing) = 0\right\}, \\ \Delta_b^z &= \left\{\mu^b \in \Delta^H(S_a) : \mu^b(\mathrm{S}|\varnothing) = 1\right\}. \end{aligned}$$

Note that Ann's initial belief and belief in the subgame must coincide.[29]

Under these belief restrictions, W *is not justifiable for Bob. Moreover,* E.L *and* E.C *are not justifiable either, because* L *and* C *are dominated by a mixture of* M *and* R *in the subgame. For Ann, strategies* N.T *and* N.U *are not justifiable under the belief restrictions. Thus, we have*

$$\begin{aligned} S^1_{a,z} &= \{\mathrm{S}, \mathrm{N.D}, \mathrm{N.B}\}, \\ S^1_{b,z} &= \{\mathrm{E.M}, \mathrm{E.R}\}. \end{aligned}$$

At step 2, Ann must assign probability 0 *to* E.L *and* E.C. *Hence, her only justifiable (reduced)*

[25] See Battigalli (1997) and Battigalli & Prestipino (2013). The main result of the latter article implies Remark **??** for all finite games.

[26] The difference between the two solution concepts in terms of epistemic foundations is analyzed in Catonini (2019).

[27] This is due to the non-monotonicity of strong belief.

[28] We conjecture that this result can be extended to the infinite games considered here.

[29] Symbols W and S respectively denote the *reduced* strategies that unilaterally terminate the game choosing W and S, that is, $\{\mathrm{W}\} \times \{\mathrm{L}, \mathrm{C}, \mathrm{M}, \mathrm{R}\}$ and $\{\mathrm{S}\} \times \{\mathrm{T}, \mathrm{U}, \mathrm{D}, \mathrm{B}\}$. To ease notation, in examples we refer to reduced strategies.

strategy is S. *For Bob, no strategy can be eliminated. So we have*

$$\begin{aligned} S^2_{a,z} &= \{\mathrm{S}\}, \\ S^2_{b,z} &= \{\mathrm{E.M}, \mathrm{E.R}\}. \end{aligned}$$

At step 3, Bob does not refine his beliefs at history $((\mathrm{N},\mathrm{E}))$*, because this history is prevented by Ann at step 2. Thus,*

$$\begin{aligned} S^\infty_{a,z} &= \{\mathrm{S}\}, \\ S^\infty_{b,z} &= \{\mathrm{E.M}, \mathrm{E.R}\}. \end{aligned}$$

Path $z = ((\mathrm{S},\mathrm{E}))$ *is credible. Note also that, in this case, all the* z*-rationalizable strategy profiles induce* z*. This means that for all beliefs that are consistent with strategic reasoning, players decide to stay on the agreed-upon path.*

To give a sense of how z-rationalizability disciplines behavior and beliefs beyond what is revealed by Example 2, we relate it to some solution concepts. To make the following observations self-contained we need some preliminary definitions.

A strategy profile s^* is a (pure) **self-confirming equilibrium** (under perfect feedback) if there is a profile of beliefs (subjective probability measures) $\left(\mu^i_\varnothing\right)_{i\in I} \in \times_{i\in I}\Delta\left(S_{-i}\left(\zeta\left(s^*\right)\right)\right)$ such that, for each player $i \in I$, $s^*_i \in \arg\max_{s_i\in S_i}\mathbb{E}_{\mu_\varnothing}\left(u_i\left(\zeta\left(s_i,\cdot\right)\right)\right)$; that is, each player i's strategy s^*_i is an ex-ante best reply to an initial belief that predicts the path of play induced by s^* and is therefore confirmed by what i observes if s^* is played, although it need not be the correct belief.[30]

A strategy profile s^* is **strongly rationalizable** if $s^* \in \times_{i\in I}S^\infty_i$, where $\left(S^n_i\right)_{i\in I, n\in\mathbb{N}}$ is the iterated deletion procedure obtained from the definition of z-rationalizability by replacing set of CPSs $\Delta^z_{i,\mathrm{sb}}(S^{n-1}_{-i,z}) = \Delta^z_i \cap \Delta^H_{\mathrm{sb}}(S^{n-1}_{-i,z})$ with $\Delta^H_{\mathrm{sb}}(S^{n-1}_{-i,z})$ for each i and n, i.e., by removing the requirement of belief in the path and keeping only the strong belief requirement.

Remark 2 *Every credible path is induced by some self-confirming equilibrium.*

Proof. Let z be credible, that is, $\times_{i\in I}S^\infty_{i,z} \neq \emptyset$. We first prove the following:

Claim $z = \zeta\left(s^*\right)$ *for some* $s^* \in \times_{i\in I}S^\infty_{i,z}$.

Proof of the claim. By way of contraposition, suppose that $z \neq \zeta\left(s^*\right)$ for all $s^* \in \times_{i\in I}S^\infty_{i,z}$, that is,

$$\emptyset = S\left(z\right) \cap \left(\times_{i\in I}S^\infty_{i,z}\right) = \times_{i\in I}\left(S_i\left(z\right) \cap S^\infty_{i,z}\right),$$

[30] For finite games, this is equivalent to the definition of pure-strategy self-confirming equilibrium of Fudenberg & Levine (1993), who maintained that the path of play is ex post observable. This special assumption is not part of the general definition of self-confirming equilibrium (see, for example, Battigalli *et al.* 2015). The best reply condition can be strengthened by requiring $s^*_i \in r_i\left(\mu^i\right)$ for a CPS $\bar{\mu}^i$ such that $\bar{\mu}^i\left(\cdot|\varnothing\right) = \mu^i_\varnothing$ without changing the paths that can be induced.

where the second equality holds because $S(z) = \times_{i \in I} S_i(z)$. Then $S_j(z) \cap S^n_{j,z} = \emptyset$ for some $j \in I$ and $n \in \mathbb{N}$. This implies that $\Delta^z_i \cap \Delta^H_{\text{sb}}(S^n_{-i,z}) = \emptyset$ and, therefore, $S^{n+1}_{i,z} = \emptyset$ for some other player $i \in I \backslash \{j\}$, so that $\times_{i \in I} S^\infty_{i,z} = \emptyset$. □

With this, pick any $s^* = (s^*_i)_{i \in I} \in \times_{i \in I} S^\infty_{i,z}$ such that $z = \zeta(s^*)$. Then there is a profile of CPSs $(\mu^i)_{i \in I} \in \times_{i \in I} \Delta^z_i$ such that $s^*_i \in r_i(\mu^i)$ for each $i \in I$. This implies that $\mu^i(S_{-i}(\zeta(s^*)) | \varnothing) = 1$ and $s^*_i \in \arg\max_{s_i \in S_i} \mathbb{E}_{\mu^i(\cdot|\varnothing)}(u_i(\zeta(s_i, \cdot)))$ for each $i \in I$, that is, s^* is a self-confirming equilibrium inducing path z. ■

Battigalli (1987) shows that every (pure) self-confirming equilibrium pair (under perfect feedback) of a two-person game is realization-equivalent to a randomized Nash equilibrium. Therefore, we obtain the following:

Remark 3 *In a two-person game, every credible path is induced with probability* 1 *by some randomized Nash equilibrium pair.*

Proof. Let z be credible. By Remark 2, there is a self-confirming equilibrium pair s^* such that $z = \zeta(s^*)$. Therefore, there is a pair of beliefs $(\mu^i_\varnothing)_{i \in I} \in \times_{i \in I} \Delta(S_{-i})$ such that $\mu^i_\varnothing(S_{-i}(z)) = 1$ and $s^*_i \in \arg\max_{s_i \in S_i} \mathbb{E}_{\mu^i_\varnothing}(u_i(\zeta(s_i, \cdot)))$ for each $i \in I$. For each i, take any behavior strategy $\sigma_i \in \times_{h \in H} \Delta(A_i(h))$ such that, for all $h \in H$ and $a_i \in A_i(h)$,

$$h \prec z \Longrightarrow \sigma_i(s^*_i(h) | h) = \mu^{-i}_\varnothing(\{s_i \in S_i(h) : s_i(h) = s^*_i(h)\}) = 1,$$

$$\left(h \not\prec z \text{ and } \mu^{-i}_\varnothing(S_i(h)) > 0\right) \Longrightarrow \sigma_i(a_i | h) = \frac{\mu^{-i}_\varnothing(s_i \in S_i(h) : s_i(h) = a_i)}{\mu^{-i}_\varnothing(S_i(h))}.$$

For histories not covered by these two cases σ_i is arbitrary. Since $\mu^{-i}_\varnothing(S_i(z)) = 1$, if $h \prec z$, $-i$ is certain that i takes the on-path action $s^*_i(h)$; if $h \not\prec z$ and $\mu^{-i}_\varnothing(S_i(h)) > 0$ it must be the case that only $-i$ deviated from the path z in history h. Therefore, for all histories such that i did not deviate from path z, the behavior strategy of i is obtained from the belief of the co-player $-i$. By construction, behavior strategy pair σ induce path z with probability 1, and s^*_{-i} is a best reply to σ_i for each i. Since the expected payoff of both s^*_{-i} and σ_{-i} is $u_{-i}(z)$, also σ_{-i} is a best reply to σ_i. Thus, σ is a randomized Nash equilibrium inducing z with probability 1. ■

Catonini (2020) shows that, in finite games, if $\times_{i \in I} S^\infty_{i,z} \neq \emptyset$ then there is a strongly rationalizable strategy profile inducing path z. Thus, we obtain the following:

Remark 4 *Fix arbitrarily a path $z \in Z$ in a finite game. If every strongly rationalizable profile induces z, then z is the unique credible path.*

To illustrate this result: in the standard Battle of the Sexes with an outside option, only the path induced by the standard forward-induction strategy pair is credible; in the finitely repeated

Prisoners Dilemma only the backward induction path is credible; and (by Battigalli 1997, Theorem 4) in every finite game with perfect information and no relevant ties only the backward induction path is credible.

3 Results

We say that a path $z \in Z$ is credible under the uncertainty attitudes (captured by) $(v_i)_{i\in I}$ when it is credible under the payoff functions $(u_i)_{i\in I} = (v_i \circ \pi_i)_{i\in I}$. Fix a profile of baseline uncertainty attitudes $(\overline{v}_i)_{i\in I}$ (e.g., the profile of identity functions representing uncertainty neutrality). We say that $(v_i)_{i\in I}$ features **more uncertainty aversion** than $(\overline{v}_i)_{i\in I}$ when each v_i is a concave transformation of $\overline{v}_i$.

Theorem 1 *Fix a path $z \in Z$ and uncertainty attitudes $(v_i)_{i\in I}, (\overline{v}_i)_{i\in I}$, where $(v_i)_{i\in I}$ features more uncertainty aversion than $(\overline{v}_i)_{i\in I}$. If z is credible under $(\overline{v}_i)_{i\in I}$, then it is credible also under $(v_i)_{i\in I}$.*

Theorem 1 claims that, as players become more uncertainty averse, more paths of the game are credible agreements. As argued in the Introduction, it is easy to find examples of paths that are not credible under low uncertainty aversion but become credible under high aversion. That a credible path under lower aversion will remain credible under higher aversion may seem obvious in light of the results of Weinstein (2016) and Battigalli *et al.* (2016): higher aversion implies more possible best replies, thus a wider range of rational reactions to a deviation (and increased aversion to the uncertainty about the reaction). However, forward-induction reasoning may destroy this argument. If there are more possible reactions to a deviation, the deviation may become profitable under some beliefs, and forward-induction reasoning about the continuation plan of the deviator may ultimately refine the set of reactions. The following examples show that, even though the set of credible paths (weakly) expands with uncertainty aversion, for some credible path z the set of z-rationalizable strategy profiles may change non-monotonically.[31]

Example 3 *Consider again the game form of Example 1. Suppose now that players are uncertainty averse. In particular, let $v_i(x) = \sqrt{x}$ $(i = a, b)$:*

A\B	W	E
N	$3,3$	$\cdot-$
S	$0,0$	$3,3$

$\longrightarrow$

A\B	L	C	M	R
T	$2,\sqrt{2}$	$2,\sqrt{5}$	$2,4$	$2,0$
U	$2,\sqrt{2}$	$2,\sqrt{5}$	$2,0$	$2,4$
D	$0,\sqrt{2}$	$4,1$	$0,4$	$0,0$
B	$4,\sqrt{2}$	$0,1$	$0,0$	$0,4$

Figure 2

[31] As in Example 2, we let S and W respectively denote the *reduced* strategies of Ann and Bob that unilaterally terminate the game.

Compared to the case of uncertainty neutrality, strategy E.C *is now optimal under some* $\mu^b \in \Delta_2$*, for instance,* $\mu^b(\text{N.T}|\,((\text{N},\text{E}))) = \mu^b(\text{N.U}|\,((\text{N},\text{E}))) = 1/2$*. Thus, we have*

$$\begin{aligned} S^1_{a,z} &= \{\text{S}, \text{N.D}, \text{N.B}\}, \\ S^1_{b,z} &= \{\text{E.C}, \text{E.M}, \text{E.R}\}. \end{aligned}$$

At step 2, Ann can assign positive probability to E.C *but not to* E.L*, therefore, for every belief that makes the deviation to* N *profitable, only* N.D *is optimal. For Bob,* E.C *is not justifiable anymore. So we have*

$$\begin{aligned} S^2_{a,z} &= \{\text{S}, \text{N.D}\}, \\ S^2_{b,z} &= \{\text{E.M}, \text{E.R}\}. \end{aligned}$$

At step 3, Ann cannot assign positive probability to E.C *anymore, therefore her only justifiable strategy is* S*. Bob, at history* $((\text{N},\text{E}))$*, must now assign probability* 1 *to* N.D*, therefore his only justifiable strategy is* E.M*. Thus,*

$$\begin{aligned} S^\infty_{a,z} &= \{\text{S}\}, \\ S^\infty_{b,z} &= \{\text{E.M}\}. \end{aligned}$$

Path $z = ((\text{S},\text{E}))$ *is credible, as under uncertainty neutrality, but the set of* z*-rationalizable strategies has shrunk.*

Surprisingly, the steps of strategic reasoning triggered by an increase in uncertainty aversion can even select reactions that are not compatible with strategic reasoning under the baseline uncertainty aversion.

Example 4 *Consider again the game form of Example 1. Suppose now that players are even more uncertainty averse than in Example 3; in particular, let* $v_i(x) = x^{\frac{1}{4}}$ $(i = a, b)$*:*

A\B	W	E
N	$\sqrt{3}, \sqrt{3}$	$\cdot -$
S	$0, 0$	$\sqrt{3}, \sqrt{3}$

$\longrightarrow$

A\B	L	C	M	R
T	$2^{\frac{1}{2}}, 2^{\frac{1}{4}}$	$2^{\frac{1}{2}}, 5^{\frac{1}{4}}$	$2^{\frac{1}{2}}, 2$	$2^{\frac{1}{2}}, 0$
U	$2^{\frac{1}{2}}, 2^{\frac{1}{4}}$	$2^{\frac{1}{2}}, 5^{\frac{1}{4}}$	$2^{\frac{1}{2}}, 0$	$2^{\frac{1}{2}}, 2$
D	$0, 2^{\frac{1}{4}}$	$2, 1$	$0, 2$	$0, 0$
B	$2, 2^{\frac{1}{4}}$	$0, 1$	$0, 0$	$0, 2$

Figure 3

Compared to the case of uncertainty neutrality, also strategy E.L *is now optimal under some* $\mu^b \in$

Δ_2, *for instance,* $\mu^b(\text{N.D}|(\text{N},\text{E})) = \mu^b(\text{N.B}|(\text{N},\text{E})) = 1/2$. *Thus, we have*

$$\begin{aligned} S^1_{a,z} &= \{\text{S}, \text{N.D}, \text{N.B}\}, \\ S^1_{b,z} &= \{\text{E.L}, \text{E.C}, \text{E.M}, \text{E.R}\}. \end{aligned}$$

Given this, at step 2, Ann does not eliminate any strategy. For Bob, as before, E.C ceases to be justifiable. So we have

$$\begin{aligned} S^2_{a,z} &= \{\text{S}, \text{N.D}, \text{N.B}\}, \\ S^2_{b,z} &= \{\text{E.L}, \text{E.M}, \text{E.R}\}. \end{aligned}$$

At step 3, Ann can assign positive probability to E.L *but not to* E.C*, therefore, for every belief that makes the deviation to* N *profitable, only* N.B *is optimal. Thus,*

$$\begin{aligned} S^3_{a,z} &= \{\text{S}, \text{N.B}\}, \\ S^3_{b,z} &= \{\text{E.L}, \text{E.M}, \text{E.R}\}. \end{aligned}$$

At step 4, Bob, at history (N,E)*, must now assign probability* 1 *to* N.B*, therefore his only justifiable strategy is* E.R*. Hence at step 5 Ann eliminates* N.B*. We have*

$$\begin{aligned} S^\infty_{a,z} &= \{\text{S}\}, \\ S^\infty_{b,z} &= \{\text{E.R}\}. \end{aligned}$$

Path $z = (\text{S},\text{E})$ *is again credible. However, the set of* z*-rationalizable strategies of Bob does not contain the one obtained under uncertainty neutrality (see Example 2), and it is even disjoint from the one obtained under intermediate uncertainty aversion (see Example 3).*

This is the big challenge for the proof of Theorem 1: when uncertainty aversion increases, what guarantees that the new set of possible reactions to a deviation is still able to deter it? Before providing an intuition for this, we present the proof of the theorem.

We start with a preliminary lemma, which is an extension to sequential optimality and justifiability of the analogous result in Battigalli *et al.* (2016).[32]

Lemma 1 *Fix* $i \in I$*, compact* $E_{-i} \subseteq S_{-i}$*,* $\overline{v}_i$ *and* v_i*, where* v_i *is a concave transformation of* $\overline{v}_i$*. If strategy* s_i *is justifiable given belief set* $\Delta^z_{i,\text{sb}}(E_{-i})$ *under* $\overline{v}_i$*, so it is under* v_i*.*

Proof. Let s_i be $\overline{v}_i$-optimal under CPS $\mu^i \in \Delta^z_{i,\text{sb}}(E_{-i})$. Then s_i maximizes expected payoff under $\overline{v}_i$ and $\mu^i(\cdot|\varnothing)$. Let $\hat{S}_{-i} = \text{supp}\mu^i(\cdot|\varnothing)$. Then, by the main result of Battigalli *et al.* (2016),

[32] See also Weinstein (2016), who considers finite games.

there is $\eta \in \Delta(\hat{S}_{-i})$ such that s_i maximizes expected payoff under η and v_i. Now, for each $h \in H(s_i)$ such that $\eta(S_{-i}(h)) = 0$ and $\eta(S_{-i}(h')) > 0$ for all $h' \prec h$, since s_i maximizes expected payoff in $S_i(h)$ under $\overline{v}_i$ and $\mu^i(\cdot|h)$, find η^h as above. Proceeding in this fashion, we obtain a collection of histories $\widetilde{H}$ and corresponding beliefs $(\eta^h)_{h\in\widetilde{H}}$. With this, it is possible to construct $\widetilde{\mu}^i \in \Delta^H(S_{-i})$ so that $\widetilde{\mu}^i(\cdot|\varnothing) = \eta$ and $\widetilde{\mu}^i(\cdot|h) = \eta^h$ for each $h \in \widetilde{H}$. By construction of the supports, $\widetilde{\mu}^i$ strongly believes E_{-i}.

Since $\mu^i(S_{-i}(z)|\varnothing) = 1$ and $S_{-i}(z)$ is closed, $\text{supp}\mu^i(\cdot|\varnothing) \subseteq S_{-i}(z)$, and hence $\widetilde{\mu}^i(S_{-i}(z)|\varnothing) = 1$. For each $h \in \widetilde{H}$, if $E \cap S_{-i}(h) \neq \emptyset$, since $\mu^i(E|h) = 1$ and E is closed, $\text{supp}\mu^i(\cdot|h) \subseteq E$, and hence $\widetilde{\mu}^i(E|h) = 1$. Thus $\widetilde{\mu}^i \in \Delta^z_{i,\text{sb}}(E_{-i})$. By construction, s_i is v_i-optimal under $\widetilde{\mu}^i$. ■

Next, similarly to Battigalli & Catonini (2026) for the comparison of Strong Directed Rationalizability under stronger and weaker restrictions on the beliefs about payoff-types, we decompose the comparison of z-rationalizability under baseline and increased uncertainty aversion into a chain of pairwise comparisons along a sequence of elimination procedures, where each procedure k carries out k steps of elimination under increased aversion and the following steps under the baseline uncertainty attitudes. Thus, procedure 0 coincides with z-rationalizability under the baseline attitudes. Compared to procedure k, procedure $k+1$ postpones the stronger requirement of justifiability under the baseline attitudes from step $k+1$ to step $k+2$, so at step $k+1$ we obtain inclusion between the two procedures by Lemma 1. This inclusion is maintained in terms of outcomes for each of the following steps of elimination, because, roughly speaking, starting from a larger set of on-path behaviors results in a larger set of justifiable behaviors along those paths, given the same uncertainty attitudes. However, differently from Battigalli & Catonini (2026), an elimination procedure may not converge in finitely many steps, because we allow for infinite horizon, hence, infinite sets of strategies.[33] Therefore, we need to use an infinite sequence of elimination procedures, where for no k we obtain the whole z-rationalizability procedure under increased aversion. To overcome this problem, we will use a sort of "diagonalization argument" to identify suitable sequences of strategies, one for each step k of procedure k, and study their limit. Relatedly, the proof contains some technical details related to compactness-continuity arguments.

Proof of Theorem 1. For each $i \in I$ and $\mu^i \in \Delta^H(S_{-i})$, let $\overline{r}_i(\mu^i)$ and $r_i(\mu^i)$ denote the sets of optimal strategies under, respectively, $\overline{v}_i$ and v_i. We introduce a sequence of elimination procedures that gradually transforms z-rationalizability under $\overline{v} = (\overline{v}_i)_{i\in I}$ into z-rationalizability under $v = (v_i)_{i\in I}$. For each $k \geq 0$, define $((\mathrm{S}^n_{k,i})_{i\in I})_{n=0}^{\infty}$ as follows. For each $n = 1, ..., k$ (provided that $k > 0$) and $i \in I$, let

$$\mathrm{S}^n_{k,i} := \left\{ s_i \in \mathrm{S}^{n-1}_{k,i} : \exists \mu^i \in \Delta^z_{i,\text{sb}}(\mathrm{S}^{n-1}_{k,-i}), s_i \in r_i(\mu^i) \right\}.$$

[33] Which typically have the power of the continuum, just like the set of all countably infinite sequences of 0's and 1's.

For each $n > k$ and $i \in I$, let

$$\mathrm{S}_{k,i}^{n} := \left\{ s_i \in \mathrm{S}_{k,i}^{n-1} : \exists \mu^i \in \Delta_{i,\mathrm{sb}}^{z}(\mathrm{S}_{k,-i}^{n-1}), s_i \in \overline{r}_i(\mu^i) \right\}.$$

Recall that $\Delta_{i,\mathrm{sb}}^{z}(E_{-i}) = \Delta_i^z \cap \Delta_{i,\mathrm{sb}}(E_{-i})$, Δ_i^z is compact, $\Delta_{i,\mathrm{sb}}(\cdot)$ preserves compactness, and r_i and $\overline{r}_i$ are upper-hemicontinuous, so that $r_i(K) = \bigcup_{\mu^i \in K} r_i\left(\mu^i\right)$ (resp. $\overline{r}_i(K)$) is compact whenever K is compact. With this, by a straightforward induction argument each set $\mathrm{S}_{k,i}^{n}$ is compact. We will repeatedly use the following result.

Claim 1 *For each $n > k$, if $s_i \in \overline{r}_i(\mu^i)$ for some $\mu^i \in \Delta_i^z$ that strongly believes $\mathrm{S}_{k,-i}^{n-1}, ..., \mathrm{S}_{k,-i}^{1}$, then $s_i \in \mathrm{S}_{k,i}^{n}$.*

Proof. Fix $\mu^i \in \Delta_i^z$ that strongly believes $\mathrm{S}_{k,-i}^{n-1}, ..., \mathrm{S}_{k,-i}^{1}$ and $s_i \in \overline{r}_i(\mu^i)$. Fix $m \in \{1, ..., n\}$ and suppose by way of induction that $s_i \in \mathrm{S}_{k,i}^{m-1}$. Since $s_i \in \overline{r}_i(\mu^i)$ for some $\mu^i \in \Delta_{i,\mathrm{sb}}^{z}(\mathrm{S}_{k,-i}^{m-1})$, by Lemma 1 $s_i \in r_i(\widetilde{\mu}^i)$ for some $\widetilde{\mu}^i \in \Delta_{i,\mathrm{sb}}^{z}(\mathrm{S}_{k,-i}^{m-1})$ as well. So, even for $m \leq k$, $s_i \in \mathrm{S}_{k,i}^{m}$. $\square$

The elimination procedure with $k = 0$ coincides with z-rationalizability under $\overline{v}$. Suppose that $\mathrm{S}_0^{\infty} \neq \emptyset$, otherwise the statement of the theorem trivially holds. For each $k > 0$, S_k^k coincides with step k of z-rationalizability under v_i. We are going to show that, for every $i \in I$ and $s_i \in \mathrm{S}_{0,i}^{\infty}$, there exists $s_i^* \in \cap_{k>0} \mathrm{S}_{k,i}^{k}$ (so that s_i^* is z-rationalizable under v_i) such that $s_i^*(h) = s_i(h)$ for all $h \in H\left(\mathrm{S}_0^{\infty}\right)$; this implies the desired result.

First, for each $k > 0$, we are going to prove by induction on $n \geq k$ the following statements:

IH1(k, n) For every $i \in I$, there exists a measurable map $\varsigma_i^n : \mathrm{S}_{k-1,i}^{n} \to \mathrm{S}_{k,i}^{n}$ such that, for each $s_i \in \mathrm{S}_{k-1,i}^{n}$, $\varsigma_i^n(s_i)(h) = s_i(h)$ for all $h \in H\left(\mathrm{S}_{k-1}^{n-1}\right)$;

IH2(k, n) For every $i \in I$, there exists a measurable map $\xi_i^n : \mathrm{S}_{k,i}^{n} \to \mathrm{S}_{k-1,i}^{n-1}$ such that, for each $s_i \in \mathrm{S}_{k,i}^{n}$, $\xi_i^n(s_i)(h) = s_i(h)$ for all $h \in H\left(\mathrm{S}_{k}^{n-1}\right)$.

The basis step ($n = k$) for IH1(k, n) comes from Lemma 1: since $\mathrm{S}_k^{k-1} = \mathrm{S}_{k-1}^{k-1}$, we get $\mathrm{S}_{k-1}^{k} \subseteq \mathrm{S}_k^k$. The basis step ($n = k$) for IH2$(k, n)$ comes from the simple observation that $\mathrm{S}_k^k \subseteq \mathrm{S}_k^{k-1} = \mathrm{S}_{k-1}^{k-1}$.

IH2$(k, n+1)$ needs to be proven first, because the proof of IH1$(k, n+1)$ uses IH2$(k, n+1)$.

Fix $i \in I$, $\bar{\mu}_i$ that strongly believes $\mathrm{S}_{k-1,-i}^{n-1}$,...,$\mathrm{S}_{k-1,-i}^{0}$, and $\bar{s}_i \in S_i$ that is a continuation best reply to $\mu_i(\cdot|h)$ at *all* $h \in H$ — this is the standard notion of (not weak) sequential optimality, we omit the formalization. For each $s_i \in \mathrm{S}_{k,i}^{n+1}$, define $\xi_i^{n+1}(s_i)$ as

$$\xi_i^{n+1}(s_i) = \left((s_i(h))_{h \in H\left(\mathrm{S}_k^n\right)}, (\bar{s}_i(h))_{h \in H \setminus H\left(\mathrm{S}_k^n\right)} \right).$$

Map ξ_i^{n+1} is continuous (and hence measurable) because it is the product of two continuous maps (identity and constant). There remains to show $\xi_i^{n+1}(s_i) \in \mathrm{S}_{k-1,i}^{n}$. Fix $\mu^i \in \Delta_{i,\mathrm{sb}}^{z}(\mathrm{S}_{k,-i}^{n})$ such that

$s_i \in \overline{r}_i(\mu^i)$. Construct $\widetilde{\mu}^i \in \Delta^z_{i,\mathrm{sb}}(\mathrm{S}^{n-1}_{k-1,-i})$ that strongly believes $\mathrm{S}^{n-1}_{k-1,-i}$,...,$\mathrm{S}^1_{k-1,-i}$ as follows:[34]

1. for each $h \in H(\mathrm{S}^n_k)$, let $\widetilde{\mu}^i(\cdot|h)$ be the pushforward of $\mu^i(\cdot|h)$ through the map $\xi^n_{-i} = \left(\xi^n_j\right)_{j\neq i}$ supplied by IH2(k,n);

2. for each $h \notin H(\mathrm{S}^n_k)$ such that $\widetilde{\mu}^i(S_{-i}(h)|h') > 0$ for some $h' \prec h$ with $h' \in H(\mathrm{S}^n_k)$, derive $\widetilde{\mu}^i(\cdot|h)$ from $\widetilde{\mu}^i(\cdot|h')$ by conditioning;

3. for each $h \notin H(\mathrm{S}^n_k)$ such that $\widetilde{\mu}^i(S_{-i}(h)|h') = 0$ for all $h' \prec h$ with $h' \in H(\mathrm{S}^n_k)$, let $\widetilde{\mu}^i(\cdot|h) = \bar{\mu}_i(\cdot|h)$.

By Claim 1, $\overline{r}_i(\widetilde{\mu}^i) \subseteq \mathrm{S}^n_{k-1,i}$. Now we show $\xi^{n+1}_i(s_i) \in \overline{r}_i(\widetilde{\mu}^i)$. Note that, for every $s_{-i} \in \xi^n_{-i}(\mathrm{S}^n_{k,-i})$, $\zeta(\xi^{n+1}_i(s_i), s_{-i}) \in \zeta(\mathrm{S}^n_k)$, because $\xi^{n+1}_i(s_i)$ and s_{-i} prescribe the same moves as $s_i \in \mathrm{S}^n_{k,i}$ and some $s_{-i} \in \left(\xi^{n+1}_{-i}\right)^{-1}(s_{-i}) \subseteq \mathrm{S}^n_{k,-i}$ within $H(\mathrm{S}^n_k)$. Thus, for every $h \in H(\xi^{n+1}_i(s_i))\backslash H(\mathrm{S}^n_k)$, $h \notin H(\xi^n_{-i}(\mathrm{S}^n_{k,-i}))$, and hence $\widetilde{\mu}^i(S_{-i}(h)|h') = 0$ for all $h' \in H(\mathrm{S}^n_k)$. So, $\widetilde{\mu}^i(\cdot|h) = \bar{\mu}_i(\cdot|h)$, and hence $\xi^{n+1}_i(s_i)$ is a continuation best reply to $\widetilde{\mu}^i(\cdot|h)$, as it prescribes the same moves as $\bar{s}_i$ from h onwards. Now fix $h \in H(\xi^{n+1}_i(s_i)) \cap H(\mathrm{S}^n_k)$ and suppose by way of induction that $\xi^{n+1}_i(s_i)$ was proven to be a continuation best reply to $\widetilde{\mu}^i(\cdot|\widetilde{h})$ at each $\widetilde{h} \prec h$ (it is vacuously true at the initial history). This implies that for every continuation best reply s'_i to $\widetilde{\mu}^i(\cdot|h)$, there exists $s''_i \in \overline{r}_i(\widetilde{\mu}^i) \cap S_i(h)$ that, under $\widetilde{\mu}^i(\cdot|h)$, induces the same distribution over terminal histories as s'_i. Since $s''_i \in \mathrm{S}^n_{k-1,i}$, by IH1$(k,n)$ there exists $\widetilde{s}'_i \in \mathrm{S}^n_{k,i} \cap S_i(h)$ such that $\widetilde{s}'_i(h') = s''_i(h')$ for every $h' \in H(\mathrm{S}^{n-1}_{k-1}) \supseteq H(\mathrm{S}^n_k)$, where the inclusion comes from IH2(k,n). The distribution over terminal histories induced by $\widetilde{s}'_i \in \mathrm{S}^n_{k,i}$ and $\mu^i(\cdot|h)$, since $\mu^i(\mathrm{S}^n_{k,-i}|h) = 1$, is concentrated on $\zeta(\mathrm{S}^n_k)$. Then, as each $s_{-i} \in \mathrm{S}^n_{k,-i}$ and $\xi^{n+1}_{-i}(s_{-i})$ prescribe the same moves within $H(\mathrm{S}^n_k)$, $\widetilde{s}'_i$ induces the same distribution with $\widetilde{\mu}^i(\cdot|h)$. By $\widetilde{s}'_i(h') = s''_i(h')$ for all $h' \in H(\mathrm{S}^n_k)$, s''_i yields the same distribution under $\mu^i(\cdot|h)$ and $\widetilde{\mu}^i(\cdot|h)$ too. But then, the distribution over terminal histories induced by the continuation best reply s'_i under $\widetilde{\mu}^i(\cdot|h)$ can also be obtained under $\mu^i(\cdot|h)$. This means that the maximal expected continuation payoff under $\mu^i(\cdot|h)$ is not lower than under $\widetilde{\mu}^i(\cdot|h)$. By the argument used for $\widetilde{s}'_i$, $s_i \in \mathrm{S}^n_{k,i}$ yields the same distribution under $\mu^i(\cdot|h)$ and $\widetilde{\mu}^i(\cdot|h)$. Then, s_i is a continuation best reply not only to $\mu^i(\cdot|h)$ but also to $\widetilde{\mu}^i(\cdot|h)$. Since $\xi^{n+1}_i(s_i)$ yields the terminal history as s_i against all $s_{-i} \in \xi^n_{-i}(\mathrm{S}^n_{k,-i})$, it is a continuation best reply to $\widetilde{\mu}^i(\cdot|h)$ too.

The proof of IH1$(k,n+1)$ is analogous, and uses, for each $j \in I$, ξ^m_j in place of ς^m_j,[35] $\mathrm{S}^n_{k,j}$ in place of $\mathrm{S}^n_{k-1,j}$, $\mathrm{S}^n_{k-1,j}$,...$\mathrm{S}^0_{k-1,j}$ in place of $\mathrm{S}^{n+1}_{k,j}$,...,$\mathrm{S}^1_{k,j}$, IH2(k,n) in place of IH1(k,n), and IH1(k,n) in place of IH2$(k,n+1)$.

[34] The array $(\widetilde{\mu}^i(\cdot|h))_{h\in H}$ inherits the chain rule from μ_i and $\bar{\mu}_i$ separately: the switch from (the pushforward of) the first to the second only happens at previously unexpected histories, where the chain rule has not bite.

[35] The index m must be adjusted to consider the different combination of steps to which IH1 refers to compared with IH2.

To conclude the proof of the theorem, fix $i \in I$ and $s_i \in \mathrm{S}_{0,i}^{\infty}$. For each $k > 0$, by IH1$(1,k)$,...,IH1(k,k), there exists $\widetilde{s}_i^k \in \mathrm{S}_{k,i}^k$ such that $\widetilde{s}_i^k(h) = s_i(h)$ for all $h \in H(\mathrm{S}_0^k) \supseteq H(\mathrm{S}_0^{\infty})$. Since S_i is compact, the sequence $(\widetilde{s}_i^k)_{k>0}$ admits a subsequence that converges to some $s_i^* \in S_i$. For all $h \in H(\mathrm{S}_0^{\infty})$, since $\widetilde{s}_i^k(h) = s_i(h)$ for all $k > 0$, $s_i^*(h) = s_i(h)$ as well. Note that, since each S_k^k coincides with step k of z-rationalizability under v, we have $\mathrm{S}_{1,i}^1 \supseteq \mathrm{S}_{2,i}^2 \supseteq ....$ So, for each $k > 0$, $\mathrm{S}_{k,i}^k$ contains $(\widetilde{s}_i^j)_{j>k}$, and since it is compact, it also contains s_i^*. Therefore, we can reach the following conclusion, which proves the theorem.

Conclusion 1 *Let $(\bar{S}_{i,z}^{\infty})_{i\in I}$ and $(S_{i,z}^{\infty})_{i\in I}$ denote, respectively, the z-rationalizable strategies under $(\overline{v}_i)_{i\in I}$ and $(v_i)_{i\in I}$. For every $i \in I$ and $s_i \in \bar{S}_{i,z}^{\infty}$, there exists $s_i^* \in S_{i,z}^{\infty}$ such that $s_i^*(h) = s_i(h)$ for all $h \in H\left(\bar{S}_z^{\infty}\right)$.* ■

Comments on the proof of Theorem 1 To provide more perspective on result and its proof, we conclude this subsection by stressing some similarities and differences between our techniques and the literature on non-monotonic strategic reasoning that precedes Battigalli & Catonini (2026). Perea (2018b) and Chen and Micali (2013) prove the order independence of, respectively, Strong Rationalizability and its iterated conditional dominance counterpart. Like us, Catonini (2020) also proves outcome-inclusions between elimination procedures with belief restrictions, focusing on games with observed actions and finite horizon. In particular, similarly to us, Perea (2018b, Theorem 3.2) compares the maximal order of elimination of unjustifiable strategies (i.e., Strong Rationalizability) with any slower order through a sequence of elimination procedures that coincide with the slow order for the first k steps and proceed at full speed thereafter. However, the pairwise comparisons are carried out in different ways: Perea uses a property of "monotonicity on reachable histories" of the strong-belief reduction operator in the absence of belief restrictions, which he proves separately; we use a richer induction hypothesis, which essentially dispenses with the need of proving such property for our elimination criterion with path-restrictions. By contrast, the workhorse lemma of Catonini (2020) directly compares two procedures of interest without decomposing the problem through a sequence of intermediate procedures. While technically more complicated, this direct approach provides additional insight on why the result holds true. In particular, in our problem, we can get an intuition of why, if a path is credible under the baseline attitudes, the additional steps of forward induction that may follow a deviation under increased uncertainty aversion cannot select reactions that make the deviation profitable. Therefore, we translate here the main idea of the proof of the workhorse lemma of Catonini (2020) for our problem.

As long as a deviation from the agreed-upon path survives both procedures, more reactions to the deviation and more continuation plans of the deviator are possible under the slow procedure. However, at some point a deviation may be ruled out under the fast procedure, but not under the slow procedure. Nonetheless, if the slow elimination procedure were to ultimately select a set of

reactions that, despite the increased uncertainty aversion, could not deter the deviation any longer, by further refining this set and the corresponding justifiable continuation plans of the deviator we would obtain a set of substrategies that is "closed under rational behavior" (in a sophisticated, dynamic sense). But then, this set would survive through (the projection in the subgame of) the fast elimination procedure as well, contradicting the initial hypothesis that the deviation is eliminated. Complications arise from the fact that the subgame that follows the deviation can be a dynamic game in itself, and this makes the problem recursive: the set of substrategies obtained from the slow procedure shall not truly survive the fast procedure as such, but an outcome-equivalent set would.

Relaxation of common knowledge of the uncertainty attitudes. Suppose now that there is no common knowledge of the exact uncertainty aversion of players. We want to study whether the set of credible paths expands as we increase the maximal levels of uncertainty aversion $(\overline{v}_i)_{i\in I}$ that players deem possible, i.e., after a concave transformation of each $\overline{v}_i$. For simplicity, we assume that the "upper bounds" $(\overline{v}_i)_{i\in I}$ are commonly known, but the argument can be extended to heterogeneous perceptions of the upper bounds, of each other's perceptions of the upper bounds, and so on. Fix an agreed-upon path $z \in Z$ and let $r_i\left(\mu^i, v_i\right)$ denote the set of weak sequential best replies to μ^i under payoff function $v_i \circ \pi_i$. With this, consider the following (decreasing) sequence $((\bar{S}^n_{i,z})_{i\in I})^\infty_{n=0}$ of subsets:

- $(n = 0)$ for each $i \in I$, let $\bar{S}^0_{i,z} = S_i$;
- $(n \in \mathbb{N})$ for each $i \in I$ and $s_i \in \bar{S}^{n-1}_{i,z}$, let $s_i \in \bar{S}^n_{i,z}$ if and only if $s_i \in r_i(\mu^i, \hat{v}_i)$ for some $\mu^i \in \Delta^z_{i,\text{sb}}(\bar{S}^{n-1}_{-i,z})$ and some convex transformation $\hat{v}_i$ of $\overline{v}_i$;
- $(n = \infty)$ for each $i \in I$, let $\bar{S}^\infty_{i,z} = \cap_{n>0}\bar{S}^n_{i,z}$.

Each set $\bar{S}^n_{i,z}$ collects the strategies of player i that, for some convex transformations of $(\overline{v}_j)_{j\neq i}$, can be justified under (the belief in path z and) strong belief in the strategies of the co-players which survived the previous step of elimination, i.e., strong belief in $\bar{S}^{n-1}_{-i,z}$. This amounts to assuming that (i) the upper bounds are "commonly known" and (ii) if players observe behavior that is inconsistent with the level of uncertainty aversion of the opponents they expected, they do not immediately deem the opponents irrational; rather, they try to rationalize their behavior with different uncertainty attitudes than previously thought, as long as this is possible below the upper bounds. Analogous considerations hold for higher-order beliefs in rationality and in each other's uncertainty attitudes.

By Lemma 1, given the set of beliefs $\Delta^z_{i,\text{sb}}(\bar{S}^{n-1}_{-i,z})$, the strategies that can be justified under some convex transformation of $\overline{v}_i$ coincide with those that can be justified precisely under $\overline{v}_i$. It follows that $((\bar{S}^n_{i,z})_{i\in I})^\infty_{n=0}$ *coincides with z-rationalizability under common knowledge of uncertainty attitudes* $(\overline{v}_i)_{i\in I}$. We record this observation as a formal result. For any profile $v = (v_i)_{i\in I}$, let

$((S^n_{i,z}(v))_{i\in I})^\infty_{n=0}$ denote the z-rationalizability deletion procedure defined in Section 2 for given uncertainty attitudes v (informally assumed to be commonly known).

Remark 5 $((\bar{S}^n_{i,z})_{i\in I})^\infty_{n=0} = ((S^n_{i,z}(\bar{v}))_{i\in I})^\infty_{n=0}$.

Now fix a profile of *true* attitudes $\hat{v} = (\hat{v}_i)_{i\in I}$ where each $\hat{v}_i$ is a *convex transformation* of $\overline{v}_i$ (thus, the upper bounds $(\overline{v}_i)_{i\in I}$ are correct). In light of the considerations above, the strategies of player i that are consistent with common knowledge of the upper bounds are those that, at each step of reasoning n, can be justified by some $\mu^i \in \Delta^z_i$ that strongly believes $\bar{S}^{n-1}_{-i,z}$ given her true attitudes $\hat{v}_i$. We say that path z is $\hat{v}$-**credible below** $(\overline{v}_i)_{i\in I}$ if, for each player i, there exists $s^*_i \in S_i(z)$ such that, for each $n > 0$, $s^*_i \in r_i(\Delta^z_{i,\text{sb}}(\bar{S}^{n-1}_{-i,z}, \hat{v}_i))$,[36] i.e., a strategy s^*_i complying with z and justified given $\hat{v}_i$ by a belief $\mu^i \in \Delta^z_i$ that strongly believes $\bar{S}^{n-1}_{-i,z}$ for each n. Contrary to what one may hope, despite the identity between z-rationalizability under common knowledge of the upper bounds and under common knowledge of the corresponding attitudes, path-monotonicity of credibility with respect to the upper bounds is not a corollary of path-monotonicity with respect to fixed attitudes, i.e., it is not a corollary of Theorem 1. This is because, by the lack of monotonicity in terms of strategies, there is no monotonicity in terms of the possible beliefs of players. Thus, the beliefs that satisfy the definition of $\hat{v}$-credibility below $(\overline{v}_i)_{i\in I}$ may not be justifiable after a concave transformation of $(\overline{v}_i)_{i\in I}$. For this reason, the proof of the next theorem will have to (re)build on the whole inductive argument of the proof of Theorem 1.

Theorem 2 *Fix a path $z \in Z$, commonly known upper bounds $v = (v_i)_{i\in I}$ and $\overline{v} = (\overline{v}_i)_{i\in I}$, where v features more uncertainty aversion than $\overline{v}$, and true uncertainty attitudes $\hat{v} = (\hat{v}_i)_{i\in I}$ below $\overline{v}$. Suppose that z is $\hat{v}$-credible below $\overline{v}$. Then, it is also $\hat{v}$-credible below v.*

Proof. For any CPS μ^i, let $\hat{r}_i(\mu^i) = r_i\left(\mu^i, \hat{v}_i\right)$ denote the set of μ^i-optimal strategies given $\hat{v}_i$. Consider z-rationalizability under common knowledge of attitudes $\overline{v}$ and v, and construct the sequence of elimination procedures of the proof of Theorem 1. For such procedures, the induction hypotheses IH1(k,n) and IH2(k,n) have been proven to hold for all k and $n \geq k$. We will use them to prove, for each k and $n \geq k$, the following statement.

Hp(k,n) For every $i \in I$, there is $\widetilde{\mu}^i \in \Delta^z_i$ that strongly believes $\mathrm{S}^n_{k,-i}, ..., \mathrm{S}^1_{k,-i}$ such that $\hat{r}_i(\widetilde{\mu}^i) \cap S_i(z) \neq \emptyset$.

For $k = 0$, for every n, Hp(k,n) holds by $\hat{v}$-credibility of z below $\overline{v}$.[37]

[36] We define the image of a set through a correspondence in the usual way: $r_i(\Delta^z_{i,\text{sb}}(\bar{S}^{n-1}_{-i,z}), \hat{\psi}_i) := \bigcup_{\mu^i \in \Delta^z_{i,\text{sb}}(\bar{S}^{n-1}_{-i,z})} r_i(\mu^i, \hat{\psi}_i)$.

[37] Recall that $((\mathrm{S}^n_{0,i})_{i\in I})^\infty_{n=0}$ coincides with z-rationalizability under $\overline{v}$. It is easy to observe that Remark 1 also holds with $\hat{r}_i(\cdot)$ in place of $r(\cdot)$.

Now fix $k > 0$ and $n \geq k$ and assume by way of induction that Hp($k-1$,n) holds. Thus, there exists $\mu^i \in \Delta_i^z$ that strongly believes $\mathrm{S}^n_{k-1,-i}$ such that $\hat{r}_i(\mu^i) \cap S_i(z) \neq \emptyset$. Let $\widetilde{\mu}^i(\cdot|\varnothing)$ be the pushforward of $\mu^i(\cdot|\varnothing)$ through the map $\varsigma^n_{-i} = \left(\varsigma^n_j\right)_{j\neq i}$ supplied by IH1(k,n), which yields $\widetilde{\mu}^i(\mathrm{S}^n_{k,-i} \cap S_{-i}(z)|\varnothing) = 1$, and complete $\widetilde{\mu}^i \in \Delta_i^z$ so that it strongly believes $\mathrm{S}^n_{k,-i}$,...,$\mathrm{S}^1_{k,-i}$. Every $\widetilde{s}'_i \in \mathrm{S}^n_{k-1,i}$ induces the same distribution over terminal histories under $\widetilde{\mu}^i(\cdot|\varnothing)$ and under $\mu^i(\cdot|\varnothing)$. For each $m \leq n$, assuming by way of induction that $\hat{r}_i(\widetilde{\mu}^i) \subseteq \mathrm{S}^m_{k,i}$, since $\hat{v}_i$ is a convex transformation of v_i (relevant for $m < k$) and $\overline{v}_i$ (relevant for $m \geq k$), for each $s_i \in \hat{r}_i(\widetilde{\mu}^i)$ Lemma 1 yields $\bar{\mu}^i \in \Delta^z_{i,\mathrm{sb}}(\mathrm{S}^m_{k,-i})$ such that $s_i \in r_i(\bar{\mu}^i, v_i)$ (for $m < k$) or $s_i \in r_i(\bar{\mu}^i, \overline{v}_i)$ (for $m \geq k$). Hence $\hat{r}_i(\widetilde{\mu}^i) \subseteq \mathrm{S}^{m+1}_{k,i}$, for each $m \leq n$; therefore, $\hat{r}_i(\widetilde{\mu}^i) \subseteq \mathrm{S}^{n+1}_{k,i}$. So, by IH2$(k,n+1)$, for every $s''_i \in \hat{r}_i(\widetilde{\mu}^i)$, there exists $\widetilde{s}'_i \in \mathrm{S}^n_{k-1,i}$ such that $\widetilde{s}'_i(h) = s''_i(h)$ for every $h \in H(\mathrm{S}^n_k) \supseteq H(\mathrm{S}^n_{k-1})$, where the inclusion comes from IH1(k,n). The distribution over terminal histories induced by $\widetilde{s}'_i \in \mathrm{S}^n_{k-1,i}$ and $\mu^i(\cdot|\varnothing)$ is concentrated on $\zeta(\mathrm{S}^n_{k-1})$, $\mu^i(\mathrm{S}^n_{k-1,-i}|\varnothing) = 1$. Then, as each $s_{-i} \in \mathrm{S}^n_{k-1,-i}$ and $\varsigma^n_{-i}(s_{-i})$ prescribe the same moves within $H\left(\mathrm{S}^n_{k-1}\right)$, $\widetilde{s}'_i$ induces the same distribution with $\widetilde{\mu}^i(\cdot|\varnothing)$. By $\widetilde{s}'_i(h) = s''_i(h)$ for all $h \in H(\mathrm{S}^n_{k-1})$, s''_i yields the same distribution under $\mu^i(\cdot|\varnothing)$ and $\widetilde{\mu}^i(\cdot|\varnothing)$ too. But then, the distribution over terminal histories induced by $s''_i \in \hat{r}_i(\widetilde{\mu}^i)$ under $\widetilde{\mu}^i(\cdot|\varnothing)$ can also be obtained under $\mu^i(\cdot|\varnothing)$. This means that the maximal expected payoff under $\mu^i(\cdot|\varnothing)$ is not lower than under $\widetilde{\mu}^i(\cdot|\varnothing)$. The same argument used for $\widetilde{\mu}^i$ shows that $\hat{r}_i(\mu^i) \subseteq \mathrm{S}^n_{k-1,i}$. Thus, every $s_i \in \hat{r}_i(\mu^i)$ yields the same expected payoff under $\mu^i(\cdot|\varnothing)$ and under $\widetilde{\mu}^i(\cdot|\varnothing)$. So, s_i is also a best reply to $\widetilde{\mu}^i(\cdot|\varnothing)$ given $\hat{v}_i$. Therefore, there exists $\widetilde{s}_i \in \hat{r}_i(\widetilde{\mu}^i)$ such that $\widetilde{s}_i \in S_i(z)$ as well. This proves Hp(k,n).

For each $i \in I$ and $k > 0$, by Hp(k,k), there exists $\widetilde{\mu}^i \in \Delta_i^z$ that strongly believes $\mathrm{S}^k_{k,-i}, ..., \mathrm{S}^1_{k,-i}$ such that $\hat{r}_i(\widetilde{\mu}^i) \cap S_i(z) \neq \emptyset$, so fix $\widetilde{s}^k_i \in \hat{r}_i(\widetilde{\mu}^i) \cap S_i(z)$. Since S_i is compact, the sequence $(\widetilde{s}^k_i)_{k>0}$ admits a subsequence that converges to some $s^*_i \in S_i$. By $\widetilde{s}^k_i \in S_i(z)$ for all $k > 0$, $s^*_i \in S_i(z)$.[38] For each $k > 0$, $(\widetilde{s}^m_i)_{m>k}$ is contained in the compact set[39]

$$\begin{aligned} & \left\{\widetilde{s}_i \in \hat{r}_i(\widetilde{\mu}^i) : \widetilde{\mu}^i \in \Delta^z_{i,\mathrm{sb}}(\mathrm{S}^k_{k,-i}) \cap ... \cap \Delta^z_{i,\mathrm{sb}}(\mathrm{S}^1_{k,-i})\right\} \\ = & \left\{\widetilde{s}_i \in \hat{r}_i(\widetilde{\mu}^i) : \widetilde{\mu}^i \in \Delta^z_{i,\mathrm{sb}}(S^k_{i,z}) \cap ... \cap \Delta^z_{i,\mathrm{sb}}(S^1_{-i,z})\right\}, \end{aligned}$$

which thus also contains the limit s^*_i. Hence, s^*_i is $\hat{v}$-credible below v. ■

Perceived uncertainty So far, in our analysis, players do not necessarily perceive strategic uncertainty: their subjective beliefs are free to assign probability 1 to a particular strategy profile of the co-players. Thus, the presence of z-rationalizable strategies that deviate from a credible path z could just be motivated by a (possibly unjustified) lack of perceived uncertainty. Therefore, we

[38] To see this, observe that, in the discounting metric, the distance between $s_i \notin S_i(z)$ and any $s'_i \in S_i(z)$ is bounded away from zero, a lower bound being determined by the finite time at which s_i deviates from z. Thus, s_i cannot be the limit of a sequence in $S_i(z)$.

[39] The argument for compactness is the same as the argument for the compactness of the z-rationalizability sets.

ask whether a minimal level of perceived uncertainty over $S^{\infty}_{-i,z}$, combined with sufficient aversion to it, makes a credible path "self-enforcing", that is, staying on path is strictly optimal for all such "cautious" beliefs.[40] For a reason we will explain soon, we focus on finite games.

Uncertainty cannot deter deviations if there is some z-rationalizable deviation that, at worst, yields the same payoff as the agreed-upon path. Thus, we rule out this knife-edge case by focusing on a credible path z under some baseline attitudes $(\overline{v}_i)_{i\in I}$ such that staying on path is not "very weakly dominated": for every $\sigma_i \in \Delta\left(S_i \backslash S_i(z)\right)$, there exists $s_{-i} \in \bar{S}^{\infty}_{-i,z} \cap S_{-i}(z)$ such that

$$\pi_i(z) > \sum_{s'_i \in S_i \backslash S_i(z)} \sigma_i(s'_i)\pi_i(\zeta(s'_i, s_{-i})). \tag{1}$$

We call a path that satisfies this condition **safely credible**. With this, for any minimal probabilities assigned to the pessimistic scenarios, a sufficient increase of uncertainty aversion would deter the deviation. In a finite game, such probabilities are guaranteed by the simple form of caution that assigns a minimal probability to every rationalizable scenario. This is the reason why we focus on finite games. So, for each $\varepsilon \in \left(0, \min_{i\in I} \frac{1}{|S_{-i}(z)|}\right)$, we say that path z is **ε-self-enforcing** if:

- $S^{\infty}_{z} \neq \emptyset$;
- for each $i \in I$ and $\mu^i \in \Delta^{z}_{i}$ that strongly believes $S^{\infty}_{-i,z}$, if $\mu^i(s_{-i}|\varnothing) \geq \varepsilon$ for each $s_{-i} \in S_{-i}(z) \cap S^{\infty}_{-i,z}$, then $r_i(\mu^i) \subseteq S_i(z)$.

We do not make any assumption about players' beliefs about the beliefs of others. In particular, we do not assume that the form of caution captured by the ε-probability requirements is commonly believed. For this reason, the definition of ε-self enforceability is not recursive.[41]

Theorem 3 *In a finite game, if a path z is safely credible under uncertainty attitudes $\overline{v} = (\overline{v}_i)_{i\in I}$, then, for every $\varepsilon \in \left(0, \min_{i\in I} \frac{1}{|S_{-i}(z)|}\right)$, there exist uncertainty attitudes $v = (v_i)_{i\in I}$ such that z is ε-self-enforcing under any uncertainty attitudes $v' = (v'_i)_{i\in I}$ that feature more uncertainty aversion than v.*

A word of caution: this is an "asymptotic" result rather than a monotonicity result, in that a path can be ε-self-enforcing for some attitudes and not after a fixed increase in uncertainty aversion, as new rationalizable reactions to deviations may kick in.

Proof. Let $(\bar{u}_i)_{i\in I} = (\overline{v}_i \circ \pi_i)_{i\in I}$. Fix any $i \in I$. Let $\widehat{z}^i$ be i's favorite path, and let $\overline{z}^i$ be i's favorite path among those that rank below z.[42] Path $\overline{z}^i$ exists by the safe credibility assumption.

[40] In Catonini (2021), self-enforceability requires the strict incentive to stay on path even in absence of perceived uncertainty. This typically requires to restrict players' beliefs with explicit threats, something that we do not do in this paper.

[41] This also avoids the "inclusion-exclusion" problem of Samuelson (1992), which would require the use of lexicographic beliefs.

[42] At least one exists, excluding the trivial case in which z is the only path of the game. If there is more than one, fix any.

Take any concave function $f_i : \mathbb{R} \to \mathbb{R}$ such that:

$$\begin{aligned} f_i(x) &= x \text{ for each } x \geq \bar{u}_i(z); \\ f_i(\bar{u}_i(\overline{z}^i)) &= \min\left\{\bar{u}_i(\overline{z}^i), \frac{\bar{u}_i(z) - (1-\varepsilon)\cdot \bar{u}_i(\widehat{z}^i)}{\varepsilon} - \varepsilon\right\}. \end{aligned}$$

Let $v_i = f_i \circ \overline{v}_i$ and $u_i = v_i \circ \pi_i = f_i \circ \bar{u}_i$ for each $i \in I$. By credibility under $\overline{v}$ and Theorem 1, z is credible under $v = (v_i)_{i\in I}$. Fix $\mu^i \in \Delta_i^{\tilde{z}}$ that strongly believes $S^\infty_{-i,z}$ such that $\mu^i(s_{-i}|\varnothing) \geq \varepsilon$ for every $s_{-i} \in S_{-i}(z) \cap S^\infty_{-i,z}$; it exists because $\varepsilon < \frac{1}{|S_{-i}(z)|}$.

Claim 2 *Under uncertainty attitudes v, for every $i \in I$ and $s_i \notin S_i(z)$, there is $s_{-i} \in S^\infty_{-i,z} \cap S_{-i}(z)$ such that $\pi_i(\zeta(s_i, s_{-i})) < \pi_i(z)$.*

Proof of the claim. Some preparations first. Let K be the number of steps needed for z-rationalizability under $(v_i)_{i\in I}$ to converge — K exists by finiteness of the game. Consider elimination procedure K in the proof of Theorem 1, which coincides with z-rationalizability under v for the first K steps and uses $\bar{v}$ thereafter. Recall that elimination procedure 0 in the proof of Theorem 1 coincides with z-rationalizability under $\bar{v}$, and fix sufficiently high n so that both procedures 0 and K converge within n steps — again, n exists by finiteness of the game. IH1(k, n) and IH2$(k, n+1)$, repeatedly applied with $k \in \{0, ..., K\}$, imply $\zeta(\mathrm{S}^n_0) = \zeta(\mathrm{S}^n_K) = \widetilde{Z}$, and in particular that ($\bigstar$) for every $s_{-i} \in \mathrm{S}^n_{0,-i}$, there is $s'_{-i} \in \mathrm{S}^n_{K,-i}$ such that $\zeta(s'_i, s_{-i}) = \zeta(s'_i, s'_{-i})$ for all $s'_i \in S_i$ with $\zeta(s'_i, s_{-i}) \in \widetilde{Z}$, and *viceversa*.

The proof is by contraposition. Suppose that there is $\hat{s}_i \in S_i \backslash S_i(z)$ such that $\pi_i(\zeta(\hat{s}_i, s_{-i})) \geq \pi_i(z)$ for all $s_{-i} \in S^\infty_{-i,z} \cap S_{-i}(z)$. Since z is safely credible, the strategies in $S_i(z)$ are best replies to some full-support probability measure ν over $\mathrm{S}^n_{0,-i} \cap S_{-i}(z)$. Observation ($\bigstar$) allows to construct a full-support probability measure.ν' over $\mathrm{S}^n_{K,-i} \cap S_{-i}(z)$ such that every $s'_i \in S_i$ with $\zeta(s'_i, s_{-i}) \in \widetilde{Z}$ for all $s_{-i} \in \mathrm{S}^n_{K,-i} \cap S_{-i}(z)$ yields the same distribution over terminal history under ν and under ν'. Also notice that every $s'_i \in S_i$ such that $\zeta(s'_i, s_{-i}) \notin \widetilde{Z}$ for some $s_{-i} \in \mathrm{S}^n_{K,-i} \cap S_{-i}(z)$ cannot be a best reply to ν'; if it were, for any CPS μ^i with $\mu^i(\cdot|\emptyset) = \nu'$, s'_i would also be a continuation best reply to $\mu^i(\cdot|h) = \nu'|S_{-i}(h)$ at every $h \prec \zeta(s'_i, s_{-i})$, which would yield $s''_i \in \mathrm{S}^n_{i,K}$ such that $\zeta(s''_i, s_{-i}) \notin \widetilde{Z}$, a contradiction. These two things combined imply that any $s'_i \in S_i(z)$ is a best reply to ν' as well. But then, given $\mathrm{S}^n_{K,-i} \subseteq S^\infty_{-i,z}$, so is $\hat{s}_i$. Thus, in turn, $\zeta(\hat{s}_i, s_{-i}) \in \widetilde{Z}$ for all $s_{-i} \in \mathrm{S}^n_{K,-i} \cap S_{-i}(z)$. By observation ($\bigstar$), this means that $\hat{s}_i$ yields the same terminal histories under $\mathrm{S}^n_{K,-i} \cap S_{-i}(z) \subseteq S^\infty_{-i,z} \cap S_{-i}(z)$.and under $\mathrm{S}^n_{0,-i} \cap S_{-i}(z) = \bar{S}^\infty_{-i,z} \cap S_{-i}(z)$. So, $\pi_i(\zeta(\hat{s}_i, s_{-i})) \geq \pi_i(z)$ for all $s_{-i} \in \bar{S}^\infty_{-i,z} \cap S_{-i}(z)$ as well. Since $\hat{s}_i \notin S_i(z)$, this contradicts safe credibility of z. $\square$

Given the Claim, for every $s_i \notin S_i(z)$ the expected payoff given $\mu^i(\cdot|\emptyset)$ is lower than $\varepsilon u_i(\overline{z}^i) +$

$(1-\varepsilon)u_i(\widehat{z}^i)$. In turn, by construction of v_i, we have

$$\varepsilon u_i(\overline{z}^i) + (1-\varepsilon)u_i(\widehat{z}^i) \leq \bar{u}_i(z) - (1-\varepsilon)\cdot\bar{u}_i(\widehat{z}^i) - \varepsilon^2 + (1-\varepsilon)u_i(\widehat{z}^i) = u_i(z) - \varepsilon^2. \tag{2}$$

Hence, the set of best replies to $\mu^i(\cdot|\varnothing)$ is contained in $S_i(z)$. So, $r_i(\mu^i) \subseteq S_i(z)$.

To conclude, notice that the argument can be repeated under any uncertainty attitudes $v' = (v'_i)_{i\in I}$ that feature more uncertainty aversion than v, as the strict inequality

$$\varepsilon u_i(\overline{z}^i) + (1-\varepsilon)u_i(\widehat{z}^i) < u_i(z)$$

implied by equation (2) is preserved by any further concave transformation of f_i. ■

4 Discussion

In this final section we discuss some extensions, modifications, and applications of our analysis.

Changes in risk and ambiguity attitudes Throughout the paper, we summarized each player i's uncertainty attitudes with a transformation v_i of her monetary payoff function π_i. Under ambiguity neutrality, v_i captures attitudes toward risk. In this section, we argue that our analysis can be extended to allow for ambiguity aversion, and we show how increases of ambiguity or risk aversion translate into the increases of "uncertainty aversion" studied by Theorems 1, 2 and 3.

We adopt the smooth-ambiguity model of choice over acts involving both uncertainty and objective risk of Klibanoff *et al.* (2005). For notational simplicity we consider two-player games. The set of states from i's perspective is the set of the co-player's pure strategies S_{-i}. If there is objective randomness, either due to random matching in a population game, or randomized behavior, player i is uncertain about the true mixed strategy $\sigma_{-i} \in \Sigma_{-i} = \Delta(S_{-i})$ and holds a subjective belief $\beta^i \in \Delta(\Sigma_{-i})$. The objective function is

$$\int_{\Sigma_{-i}} \phi_i \left(\sum_{s_{-i}\in S_{-i}} \sigma_{-i}(s_{-i})\, v_i(\pi_i(\zeta(s_i, s_{-i}))) \right) \beta^i(\mathrm{d}\sigma_{-i}),$$

where the "second-order utility function" ϕ_i is a concave (under weak ambiguity aversion) and strictly increasing function representing attitudes toward ambiguity. Under ambiguity neutrality (i.e., subjective expected utility maximization) ϕ_i is linear and only the predictive probabilities $\mu^i(s_{-i}) = \int_{\Sigma_{-i}} \sigma_{-i}(s_i)\,\beta^i(\mathrm{d}\sigma_{-i})$ matter. Battigalli *et al.* (2016) show that in simultaneous-move games an increase in ambiguity aversion by means of a concave and strictly increasing transformation of ϕ_i yields a larger set of justifiable strategies (see also Battigalli et al. 2015). As for sequential games, we assume that belief β^i is updated by Bayes rule whenever possible, as we take the correct

use of the rules of conditional probability as a cognitive rationality assumption (cf. Battigalli *et al.* 2019). It is known that in sequential decision problems and games this decision criterion is not dynamically consistent. This lack of dynamic consistency prevents the extension of the aforementioned monotonicity result from simultaneous-move games to sequential games. Yet, we consider an environment where objective randomness is absent: players are not drawn at random from large populations, there are no chance moves, and players cannot *irreversibly* delegate their choices to randomization devices. Therefore, they only play pure strategies.[43] Assuming that this is transparent, beliefs β^i can only assign positive probability to Dirac measures $\delta_{s_{-i}}$ and are therefore equivalent to predictive beliefs $\mu^i \in \Delta(S_{-i})$. In finite games, the objective function becomes

$$\sum_{s_{-i} \in S_{-i}} \beta^i \left(\delta_{s_{-i}}\right) \phi_i \left(v_i \left(\pi_i \left(\zeta \left(s_i, s_{-i}\right)\right)\right)\right) = \sum_{s_{-i} \in S_{-i}} \mu^i \left(s_{-i}\right) \left(\phi_i \circ v_i\right) \left(\pi_i \left(\zeta \left(s_i, s_{-i}\right)\right)\right)$$

(also considering updated/revised beliefs conditional on observed histories), and dynamic consistency holds. An analogous equality holds for infinite games. Composite function $\psi_i = \phi_i \circ v_i$ captures uncertainty aversion in the assumed environment. Let τ be a concave and strictly increasing transformation; since $\tau \circ \psi_i = (\tau \circ \phi_i) \circ v_i$, a τ-concavification of ψ_i is equivalent to a τ-concavification of the second-order utility function ϕ_i, that is, to an increase in ambiguity aversion. If ϕ_i is the identity, i.e., without loss under ambiguity neutrality, $\tau \circ \psi_i = \tau \circ v_i$, and a τ-concavification is also equivalent to an increase in risk aversion. If ϕ_i is not linear, a τ-concavification need not yield a concavification of ψ_i, that is, an increase in risk aversion does not imply an increase in uncertainty aversion. Once we understand an increase in uncertainty aversion as an increase in ambiguity aversion, or in risk aversion under ambiguity neutrality, the analysis of Section 3 applies *verbatim* with v_i replaced by ψ_i.

Alternative notions of credibility Our notion of credibility is based on the assumption that beliefs are disciplined by forward-induction reasoning. Specifically, we assume Rationality and Common *Strong* Belief in Rationality under transparency of expected compliance with the path-agreement. One could consider alternative assumptions about strategic reasoning, such as Rationality and Common *Initial* Belief in Rationality under transparency of expected compliance with the path-agreement. We do not endorse this weaker notion of credibility, because according to it players are not assumed to interpret deviations as purposeful choices made to increase payoff given strategically sophisticated beliefs. Be that as it may, we note that, with it, monotonicity of the set of credible path-agreements easily follows from Battigalli *et al.* (2016) and an induction argument, for the simple reason that – unlike strong belief – initial belief is monotone.[44] The same holds for

[43] Players could plan to make their choices depend on the realizations of extraneous random variables. But such plans have to be incentive compatible, implying that for every realization the planned choice must be optimal among the pure choices. Generically, there is only one such choice.

[44] Furthermore, statements analogous to Remarks 2 and 3 hold also for this weaker version of credibility.

other forms of strategic reasoning that rely on monotone belief operators.

Time and uncertainty A sufficient condition to make sense of how we analyze uncertainty aversion is that the given game Γ with monetary payoffs is a one-period game, with payoffs accruing to the players at the end of the period. Our analysis can be extended to multi-period games, e.g., repeated games, under some assumptions about the interplay of time and uncertainty preferences. Specifically, consider an infinite-horizon game with one-period payoffs in a bounded interval $B \subseteq \mathbb{R}$. Suppose that i's preferences $\succsim_i$ over *deterministic* streams of payoffs $\mathbf{m} = (m_t)_{t\in\mathbb{N}} \in B^{\mathbb{N}}$ satisfy (time-invariant)[45] exponential discounting, with discount factor $\gamma_i \in (0,1)$:

$$\mathbf{m}' \succsim_i \mathbf{m}'' \iff \sum_{t\in\mathbb{N}} \gamma_i^{t-1} \bar{v}_i\left(m_t'\right) \geq \sum_{t\in\mathbb{N}} \gamma_i^{t-1} \bar{v}_i\left(m_t''\right),$$

where $\bar{v}_i$ is the one-period utility of money. Given subjective belief $\mu_i \in \Delta\left(B^{\mathbb{N}}\right)$ about *uncertain* streams of payoffs, we assume that the objective function is

$$\int_{B^{\mathbb{N}}} \varphi_i \left((1-\gamma_i) \sum_{t\in\mathbb{N}} \gamma_i^{t-1} \bar{v}_i m_t \right) \mu_i\left(\mathrm{d}\mathbf{m}\right)$$

where φ_i is concave and strictly increasing (cf. Dillenberger *et al.* 2024, Assumptions 1 and 2). With this, increased uncertainty aversion corresponds to a concave and strictly increasing transformation of φ_i, and our analysis applies.

Uncertainty aversion and collusion Since collusion is illegal, collusive agreements are necessarily non-binding. We argue that non-binding agreements about what to do rather than on how to react to deviations are worth exploring, because there may be little reason to have faith in the agreement after it has been violated and players may want to refrain to propose prescriptions about such reactions. Therefore, we suggest that non-binding agreements in sequential games leave room for the perception of uncertainty about co-players' off-path behavior, although perceived uncertainty is disciplined by strategic reasoning. With this, our analysis suggests the collusion is easier when uncertainty aversion is higher. We argued above that our analysis also applies to infinitely repeated games,[46] implying that cartels should find it easier to support high prices when uncertainty aversion is high. Ownership and governance changes may affect this channel by altering decision-makers' effective exposure to uncertain profits or their attitudes toward uncertainty. An IPO is a possible example, but its implications depend on the accompanying changes in ownership,

[45] For the sake of simplicity.

[46] As long as the one-period game being repeated is finite. Also, our analysis does not cover the case of random termination of the game, as this would yield also on-path uncertainty. But it is reasonable to assume that our conclusions also hold when the probability of termination is commonly believed to be very small.

control, incentives, and information.

References

[1] Anderson, L., B. Freeborn, and J. Hulbert (2012): "Risk Aversion and Tacit Collusion in a Bertrand Duopoly Experiment," *Review of Industrial Organization*, 40, 37-50.

[2] Banks J.S. and J. Sobel (1987): "Equilibrium Selection in Signaling Games," *Econometrica,* 55, 647-661.

[3] Battigalli P. (1987): *Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali*, unpublished thesis, Bocconi University.

[4] Battigalli, P. (1997): "On Rationalizability in Extensive Games," *Journal of Economic Theory*, 74, 40-61.

[5] Battigalli P. (2003): "Rationalizability in Infinite, Dynamic Games of Incomplete Information," *Research in Economics,* 57, 1-38.

[6] Battigalli, P., and E. Catonini (2024): "The Epistemic Spirit of Divinity," *Journal of Economic Theory*, 222, 105903.

[7] Battigalli, P., and E. Catonini (2026): "Monotonicity and Robust Implementation Under Forward-Induction Reasoning," *Econometrica*, 94, 505–536.

[8] Battigalli, P., and A. Prestipino (2013): "Transparent Restrictions on Beliefs and Forward Induction Reasoning in Games with Asymmetric Information," *The B.E. Journal of Theoretical Economics (Contributions)*, 13, 1-53.

[9] Battigalli, P., and M. Siniscalchi (1999): "Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic Games," *Journal of Economic Theory*, 88, 188-230.

[10] Battigalli P. and M. Siniscalchi (2002): "Strong Belief and Forward Induction Reasoning," *Journal of Economic Theory*, 106, 356-391.

[11] Battigalli P. and M. Siniscalchi (2003): "Rationalization and Incomplete Information," *The B.E. Journal of Theoretical Economics*, 3, 1-46.

[12] Battigalli, P., and M. Siniscalchi (2007): "Interactive Epistemology in Games with Payoff Uncertainty," *Research in Economics*, 61, 165-184.

[13] Battigalli, P., and P. Tebaldi (2019): "Interactive Epistemology in Simple Dynamic Games with a Continuum of Strategies," *Economic Theory*, 68, 737-763.

[14] Battigalli, P., E. Catonini, and N. De Vito (2025, work in progress): *Game Theory: Analysis of Strategic Thinking*. Bocconi University.

[15] Battigalli, P., E. Catonini, and J. Manili (2023): "Belief Change, Rationality, and Strategic Reasoning in Sequential Games," *Games and Economic Behavior*, 142, 527-551.

[16] Battigalli, P., E. Catonini, G. Lanzani, and M. Marinacci (2019): "Ambiguity Attitudes and Self-Confirming Equilibrium in Sequential Games," *Games and Economic Behavior*, 115, 1-29.

[17] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2015): "Self-Confirming Equilibrium and Model Uncertainty," *American Economic Review*, 105, 646-677.

[18] Battigalli, P., S. Cerreia-Vioglio, F. Maccheroni, and M. Marinacci (2016): "A Note on Comparative Ambiguity Aversion and Justifiability," *Econometrica*, 84, 1903-1916.

[19] Ben-Porath, E. (1997): "Rationality, Nash Equilibrium, and Backward Induction in Perfect Information Games," *Review of Economic Studies*, 64, 23-46.

[20] Bernhardt, D., and M. Rastad (2016): "Collusion under Risk Aversion and Fixed Costs," *Journal of Industrial Economics*, 64, 808-834.

[21] Catonini, E., (2019): "Rationalizability and Epistemic Priority Orderings," *Games and Economic Behavior*, 114, 101-117.

[22] Catonini, E., (2020): "On Non-monotonic Strategic Reasoning," *Games and Economic Behavior*, 120, 209-224.

[23] Catonini, E., (2021): "Self-enforcing Agreements and Forward Induction Reasoning," *Review of Economic Studies,* 88, 610-642.

[24] Chen, J., and S. Micali (2013): "The Order Independence of Iterated Dominance in Extensive Games," *Theoretical Economics*, 8, 125-163.

[25] Cho, I.K., and D. Kreps (1987): "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 102, 179-221.

[26] Dillenberger, D., D. Gottlieb, and P. Ortoleva (2024): "Stochastic Impatience and the Separation of Time and Risk Preferences," mimeo.

[27] Eichberger, J., and D. Kelsey (2011) "Are the Treasures of Game Theory Ambiguous?" *Economic Theory*, 48, 313-339.

[28] Eichberger J., S. Grant, and D. Kelsey (2018): "Ambiguity and the Centipede Game: Strategic Uncertainty in Multi-Stage Games with Almost Perfect Information," mimeo.

[29] Eichberger, J., D. Kelsey, and B. Schipper (2009): "Ambiguity and Social Interaction," *Oxford Economic Papers*, 61,355-379.

[30] Friedenberg, A. (2019): "Bargaining Under Strategic Uncertainty: The Role of Second-Order Optimism," *Econometrica*, 87, 1835-1865.

[31] Fudenberg, D., and D.K. Levine (1993): "Self-Confirming Equilibrium," *Econometrica,* 61, 523-545.

[32] Gilboa, I., and D. Schmeidler (1989): "Maxmin Expected Utility with a Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141-153.

[33] Govindan, S. and R. Wilson (2009): "On Forward Induction," *Econometrica*, 77, 1-27.

[34] Greenberg, J. (2000): "The Right to Remain Silent," *Theory and Decision*, 48, 193–204.

[35] Greenberg, J., S. Gupta, and X. Luo (2009): "Mutually Acceptable Courses of Action," *Economic Theory*, 40, 91-112.

[36] Hanany, E., P. Klibanoff, and S. Mukerji (2020): "Incomplete Information Games with Ambiguity Averse Players," *American Economic Journal: Microeconomics*, 12(2), 135-187.

[37] Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A Smooth Model of Decision Making Under Ambiguity," *Econometrica*, 73, 1849-1892.

[38] Kohlberg, E. and J.F. Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54(5), 1003-1037.

[39] Kuhn, H.W. (1953): "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games II*, ed. by H.W. Kuhn and A.W. Tucker. Princeton: Princeton University Press, 193-216.

[40] Lo, C.K. (1999): "Extensive Form Games with Uncertainty Averse Players," *Games and Economic Behavior*, 28, 256–270.

[41] Osborne, M., and A. Rubinstein (1994): *A Course in Game Theory.* Cambridge MA: MIT Press.

[42] Pearce, D. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection", *Econometrica*, 52, 1029-1050.

[43] Perea, A. (2014): "Belief in the Opponents' Future Rationality," *Games and Economic Behavior*, 83, 231-254.

[44] Perea, A. (2018a): "Order Independence in Dynamic Games," working paper.

[45] Perea, A. (2018b): "Why Forward Induction leads to the Backward Induction Outcome: A New Proof for Battigalli's Theorem", *Games and Economic Behavior*, 110, 120–138.

[46] Perea, A. (2025): "More Reasoning, Less Outcomes: A Monotonicity Result for Reasoning in Dynamic Games," EPICENTER W.P. 32, Maastricht University.

[47] Renyi, A. (1955): "On a New Axiomatic Theory of Probability," *Acta Mathematica Academiae Scientiarum Hungaricae,* 6, 285-335.

[48] Riedel, F., and L. Sass (2014): "Ellsberg Games," *Theory and Decision,* 76, 469-509.

[49] Rojas, M, and D. Vergara (2021): "Ambiguity and long-run cooperation in strategic games," *Oxford Economic Papers*, 73, 1077-1098.

[50] Rubinstein, A. (1991): "Comments on the Interpretation of Game Theory," *Econometrica*, 59, 909-924.

[51] Sabater-Grande, G., and N. Georgantzis (2002): "Accounting for Risk Aversion in Repeated Prisoners' Dilemma Games: An Experimental Test," *Journal of Economic Behavior and Organization*, 48, 37–50.

[52] Samuelson, L. (1992): "Dominated Strategies and Common Knowledge," *Games and Economic Behavior*, 4, 284-313.

[53] Schmeidler, D. (1989): "Subjective Probability and Expected Utility without Additivity," *Econometrica*, 57, 571-587.

[54] Sobel, J., L. Stole, and I. Zapater (1990): "Fixed-Equilibrium Rationalizability in Signaling Games," *Journal of Economic Theory*, 52, 304-331.

[55] Weinstein, J. (2016): "The Effect of Changes in Risk Attitude on Strategic Behavior," *Econometrica*, 84, 1881-1902.