Lexicographic Numbers and Stability in Extensive Form Games

Francesc Dilmé*

March 21, 2022

Abstract

This paper introduces the space of (lexicographic) ℓ-numbers and uses it to analyze extensive form games. We provide a simple characterization of sequential equilibria which does not use sequences of strategy profiles. We use ℓ-numbers to introduce a new concept of stability, which we call ℓ-stability, and prove the existence of ℓ-stable outcomes for all extensive form games. We show that stable outcomes are ℓ-stable when they exist and that, when there is a unique ℓ-stable outcome, it is the unique stable outcome. Together, our results provide a simple way to find stable outcomes in practice.

Key words: Lexicographic numbers, sequential equilibrium, stable outcome.

JEL classification codes: C72, C73.

1 Introduction

Selection criteria and refinements are at the core of the analysis of dynamic games with incomplete information. Indeed, many Nash equilibria in most dynamic games of economic interest fail to satisfy basic plausibility requirements, such as sequential rationality, belief consistency, or robustness to small payoff perturbations or trembles.

A number of different refinements have been suggested to make predictions more sensible. An important class consists of the limit-based equilibrium concepts, which are limits of equilibria of versions of the game where strategies are perturbed. Sequential equilibria, perfect equilibria, trembling-hand equilibria, and strategically-stable sets

*University of Bonn. Previous versions of this paper were titled "Likelihood Assignments in Extensive Form Games." fdilme@uni-bonn.de.

of equilibria are the most prominent examples. A large literature shows that these equilibrium concepts are theoretically appealing, as they satisfy the aforementioned and other plausibility requirements (e.g., backwards or forward induction).

In practice, limit-based equilibrium concepts are difficult to use: proving or disproving a given candidate satisfies some refinement typically requires finding both the right perturbation sequences and limits of corresponding equilibrium strategies, which is often not feasible in most games. Instead, weaker and less theoretically appealing selection criteria are often used, such as different versions perfect Bayesian equilibrium, iterated deletion of weakly dominated strategies, or criteria specific to signaling games (such as intuitive criterion, D1, divinity, etc).

This paper has three main contributions, aimed at making limit-based equilibrium concepts easier to use. The first contribution is to construct the space of *ℓ-numbers* to represent and simply work with asymptotic likelihoods without the need of using sequences. The second contribution is to provide a simple characterization of sequential equilibria in terms of ℓ-numbers. The last contribution is to use the language of ℓ-numbers to define *ℓ-stable outcomes*, show they always exist, and prove that satisfy forward induction and iterated strict dominance. We provide a procedure to obtain stable outcomes in practice. In a companion paper (Dilmé, 2021), we illustrate the usefulness of our approach.

We begin by defining the space of (lexicographic) *ℓ-numbers*, also referred to as *likelihoods*, which consists of the set $L \equiv \mathbb{R}_+ \times \mathbb{R}_{++}$, equipped with (a) the lexicographic order and (b) some simple operations (sum, multiplication, and division). Each ℓ-number $(\ell^{\mathrm{r}}, \ell^{\mathrm{p}}) \in L$ can be interpreted to represent the class of sequences that converge to 0 as $(n^{-\ell^{\mathrm{r}}}\, \ell^{\mathrm{p}})_n$. Hence, each ℓ-number of the form $(0, \ell^{\mathrm{p}})$ is identified with the real number $\ell^{\mathrm{p}} \in \mathbb{R}_{++}$. An ℓ-number $(\ell^{\mathrm{r}}, \ell^{\mathrm{p}})$ with $\ell^{\mathrm{r}} > 0$ is referred to as an *infinitesimal*, and it is used to distinguish the likelihood of events with vanishing probability.[1]

We use the language of ℓ-numbers to generalize the concept of strategy profile. An *ℓ-strategy profile* assigns a likelihood to each action, with the condition that the sum of the likelihoods of the actions available at each information set is 1. Assigning likelihoods to both histories and information sets is then straightforward, and done through the usual multiplication and addition of likelihoods of the actions and histories they are composed of. ℓ-Numbers permit a simple characterization of belief consistency: an assessment is consistent if and only if it is generated by some ℓ-strategy profile. An implication is that all consistent assessments can be generated by simple sequences of strategy profiles, where each action a played with zero probability in the limit is played with probability $(n^{-\lambda^{\mathrm{r}}(a)}\, \lambda^{\mathrm{p}}(a))_n$ along the sequence, for some $(\lambda^{\mathrm{r}}(a), \lambda^{\mathrm{p}}(a)) \in L$.

[1] Our construction can be interpreted as a significanlty simplified version of hyperreal numbers (see Robinson, 2016, for a review).

We define an ℓ-*equilibrium* as an ℓ-strategy profile where the likelihood of an action is not infinitesimal only if such an action is sequentially optimal. We show that each ℓ-equilibrium generates a sequential equilibrium, and that each sequential equilibrium is generated by some ℓ-equilibrium. This characterization eases obtaining sequential equilibria, as it permits working directly "in the limit" both to determine sequential optimality and belief consistency.

With the goal of studying stability in mind, we generalize the concept of tremble. An ℓ-*tremble* assigns an infinitesimal to each action, and is interpreted as the likelihood with which players make mistakes under a certain tremble. Paralleling trembling-hand perfection, an ℓ-*equilibrium for a given* ℓ-*tremble* is defined as an ℓ-strategy profile that assigns a higher likelihood than the ℓ-tremble only to actions which are sequentially-optimal. We show that ℓ-equilibria for any given ℓ-tremble can be characterized as limits of ε-perfect (instead of Nash) equilibria of perturbed versions of the game.

In the final part of our analysis, we define and characterize stability using ℓ-numbers. In the spirit of strategic stability, we say that an outcome is ℓ-*stable* if, for any ℓ-tremble, there is some ℓ-equilibrium for the ℓ-tremble with the given outcome. Remarkably, unlike stable outcomes, ℓ-stable outcomes always exist, even when payoffs are not generic. In fact, ℓ-stability is shown to be more permissive than stability: when a stable outcome exists, it is ℓ-stable. The reason is that an outcome is ℓ-stable if and only if, for any tremble, there is a strategy profile where behavior at any decision point in the game is optimal *approximately* optimal (instead of exactly optimal). We show that ℓ-stable outcomes satisfy strict versions of three important properties of stable outcomes: forward induction, iterated dominance, and invariance to reordering simultaneous moves.

We finally provide a procedure to obtain stable outcomes in practice. We first show that, when there is a unique ℓ-stable outcome, such outcome is stable as well. Then, to prove that a given outcome is the unique stable outcome, it suffices to find an ℓ-tremble with a unique ℓ-equilibrium. We illustrate the procedure through some examples.

1.1 Related Literature

Since the definition of Nash equilibrium (Nash, 1951), many refinements and selection criteria have been developed, especially for extensive form games. Perfect equilibrium (Selten, 1975), sequential equilibrium (Kreps and Wilson, 1982) or stable sets of equilibria (Kohlberg and Mertens, 1986), were celebrated for their conceptual simplicity, universal applicability, and theoretical properties. The difficulty of using them (i.e., proving that a given strategy profile satisfies or fails their conditions) have limited their use in the analysis of concrete games in applications. In practice, less theoretically appealing concepts are used, such as versions of perfect Bayesian equilibrium (Fudenberg and Tirole, 1991b) or, for signaling games, selection criteria such as the intuitive criterion and D1 (Cho and Kreps, 1987) or universal divinity (Banks and Sobel, 1987). The crucial difference between ℓ-stability and stability is that ℓ-stability requires ε-optimality instead of exact optimality along the perturbation sequence, hence it is less restrictive and ensures existence for all games. We show that, nonetheless, ℓ-stability can be used to show stability when a unique ℓ-stable outcome exists.[2]

Our paper is part of a literature that uses a different approach to study limit-base equilibrium concepts: characterizing them without explicitly working with sequences. Most saliently, conditional probability systems (CPSs, see Battigalli, 1996) and lexicographic probability systems (LPSs, see Blume, Brandenburger, and Dekel, 1991 and Govindan and Klumpp, 2003), have been used to model each player's belief about the strategies chosen by the rest of the players in perfect and proper equilibria.[34] As we discuss in detail in Appendix B, using ℓ-numbers simplifies these approaches by retaining only the largest terms in a CPS or LPS. Studying only the largest terms significantly lowers the dimensionality of the objects and suffices to characterize consistency of assessments, to evaluate sequential rationality, and to define of stable outcomes. Furthermore, as we will see, ℓ-strategy profiles and ℓ-equilibria do not require global conditions (such as the "independence property" required for LPSs or CPSs), and can be defined information set by information set.

The rest of the paper is organized as follows: Section 2 provides the notation for extensive form games used in the rest of the paper and introduces the space of ℓ-numbers. Section 3 defines ℓ-strategy profiles and ℓ-equilibria, and uses them to characterize consistent assessments and sequential equilibria. In Section 4, we define ℓ-trembles and ℓ-equilibria for a given ℓ-tremble, and characterize them in terms of limits of ε-perfect equilibria. Section 5 defines ℓ-stable outcomes and analyzes their connection to stable outcomes. Finally, Section 6 provides some examples and Section 7 concludes. Appendix A contains the omitted proofs, while Appendix B contains an exposition of the relationship between ℓ-strategy profiles, LPSs and CPSs.

---

[2]Our analysis does not consider payoff uncertainty, studied in Fudenberg, Kreps, and Levine (1988). Recently, Takahashi and Tercieux (2020) have shown that existence of outcomes robust to payoff uncertainty for generic payoffs. We show that ℓ-stable outcomes are robust to payoff perturbations.

[3]Our construction does not offer a characterization of perfect or proper equilibria. Nevertheless, recall that equilibrium paths of perfect and sequential equilibria coincidence for generic payoffs.

[4]Myerson (1986) introduces conditional probability systems to specify conditional probabilities (or beliefs) on zero-probability events, characterizes them as limits of sequences of probability distributions and uses them to characterize sequential communication equilibria and predominant communication equilibria (see also Kohlberg and Reny, 1997). Blume, Brandenburger, and Dekel (1991) employ LPSs to provide a decision-theoretic representation of preferences under lexicographic beliefs, and use it to characterize (normal-form) perfect and proper equilibria. Similarly, Mailath, Samuelson, and Swinkels (1997) use LPSs to characterize strategic independence respecting equilibria (SIRE) and compare them to proper equilibria.

2 Extensive form games and ℓ-numbers

2.1 Extensive form games

We base our definition of extensive form game on Osborne and Rubinstein (1994), although we do not follow it exactly.

Definition 2.1. A (finite) *extensive form game* $G = \langle A, H, \mathcal{I}, N, u, \rho \rangle$ has the following components.

1. **Actions and Histories:** A finite *set of actions* A and a finite set of sequences H (or *histories*) of actions satisfying:
   (a) The empty sequence $\emptyset$ is a member of H.
   (b) If $(a_j)_{j=1}^{J} \in H$ and $J' < J$ then $(a_j)_{j=1}^{J'} \in H$.

   A history $(a_j)_{j=1}^{J} \in H$ is *terminal* if there is no a_{J+1} such that $(a_j)_{j=1}^{J+1} \in H$. The set of terminal histories is denoted T. For any $h \in H$, we use $A^h \equiv \{a | (h, a) \in H\}$ to denote the set of actions available at h (notice that $A^t = \emptyset$ for all $t \in T$.) We further assume that $\cup_{h \in H} A^h = A$.

2. **Information:** A partition $\mathcal{I}$ of the non-terminal histories, such that for each *information set* $I \in \mathcal{I}$, $h \in I$, and $a \in A^h$, (a) $a \in A^{h'}$ if and only if $h' \in I$; and (b) $h'' \notin I$ for each $h'' > h$.[5] We use A^I to denote the actions available at histories in I, and I^a is the (unique) information set such that $a \in A^{I^a}$.

3. **Players:** A finite set of *players* $N \not\ni 0$ and a function $\iota : \mathcal{I} \to \{0\} \cup N$ assigning a player to each information set.
   (a) **Nature:** Player 0 is called *nature*. The function $\rho : \cup_{I \in \iota^{-1}(0)} A^I \to (0, 1]$ is such that, for each $I \in \iota^{-1}(0)$, we have $\sum_{a \in A^I} \rho(a) = 1$.
   (b) **Payoffs:** Each player $i \in N$ has an associated *(von Neumann-Morgenstern) payoff function* $u_i : T \to \mathbb{R}$. For convenience, $u_0(t) = 0$ for all $t \in T$.
   (c) **Recall:** If $(a_j)_{j=1}^{J} \in I$ and $(a_j)_{j=1}^{K} \in I'$ for some $K < J$ and $\iota(I) = \iota(I')$ then, for any $(a'_j)_{j=1}^{J'} \in I$, there is some $K' < J'$ such that $(a'_j)_{j=1}^{K'} \in I'$.

A *strategy (profile)* σ maps each action $a \in A$ into a probability $\sigma(a) \in [0, 1]$ satisfying that $\sum_{a \in A^I} \sigma(a) = 1$ for all $I \in \mathcal{I}$, and $\sigma(a) = \rho(a)$ for all a with $I^a \in \iota^{-1}(0)$ (note that a denotes a generic action in the game, not an action profile). We let Σ be the set of strategy profiles.

[5]We assume, without loss of generality, that each action only belongs to a unique information set (otherwise one can rename actions). We use $h'' > h$ to indicate that h'' succedes h.

2.2 ℓ-Numbers

In this section, we introduce the space of lexicographic (ℓ-)numbers, or likelihoods, which will be the basis of our analysis. We discuss their relationship with lexicographic probability systems (LPSs) and conditional probability systems (CPSs) in Appendix B.

We will use ℓ-numbers to represent asymptotic likelihoods with which actions are played along some sequence of strategy profiles. It will turn out that, for our analysis, it will be without loss of generality to focus on *simple sequences* of the form $(n^{-\ell^{\mathrm{r}}}\,\ell^{\mathrm{p}})_n$, for some $\ell^{\mathrm{r}} \geq 0$ indicating the *rate* at which the sequence tends to 0, and some $\ell^{\mathrm{p}} > 0$ indicating a *proportional* factor. An ℓ-number $\ell \equiv (\ell^{\mathrm{r}}, \ell^{\mathrm{p}})$ will then represent the (equivalence) class of sequences that tend to 0 as $(n^{-\ell^{\mathrm{r}}}\,\ell^{\mathrm{p}})_n$; that is,

$$(x_n)_n \in \ell \quad \Leftrightarrow \quad \lim_{n\to\infty} \frac{x_n}{n^{-\ell^{\mathrm{r}}}\,\ell^{\mathrm{p}}} = 1 \ .$$

We now define simple operations on such equivalent classes.

Definition 2.2. The *space of ℓ-numbers* is $L \equiv \mathbb{R}_+ \times \mathbb{R}_{++}$ endowed with the following order and operations for each pair $\ell \equiv (\ell^{\mathrm{r}}, \ell^{\mathrm{p}}), \hat{\ell} \equiv (\hat{\ell}^{\mathrm{r}}, \hat{\ell}^{\mathrm{p}}) \in L$:

1. *Order:* $\ell > \hat{\ell}$ if either $\ell^{\mathrm{r}} < \hat{\ell}^{\mathrm{r}}$ or $\ell^{\mathrm{r}} = \hat{\ell}^{\mathrm{r}}$ and $\ell^{\mathrm{p}} > \hat{\ell}^{\mathrm{p}}$.
2. *Addition:* $\ell + \hat{\ell} \equiv \big(\min\{\ell^{\mathrm{r}}, \hat{\ell}^{\mathrm{r}}\}\,,\, \mathbb{I}_{\ell^{\mathrm{r}} \leq \hat{\ell}^{\mathrm{r}}}\,\ell^{\mathrm{p}} + \mathbb{I}_{\ell^{\mathrm{r}} \geq \hat{\ell}^{\mathrm{r}}}\,\hat{\ell}^{\mathrm{p}}\,\big)$.
3. *Multiplication:* $\ell\,\hat{\ell} \equiv \big(\ell^{\mathrm{r}} + \hat{\ell}^{\mathrm{r}}, \ell^{\mathrm{p}}\,\hat{\ell}^{\mathrm{p}}\big)$.
4. *Division:* $\ell/\hat{\ell} \equiv \big(\ell^{\mathrm{r}} - \hat{\ell}^{\mathrm{r}}, \ell^{\mathrm{p}}/\hat{\ell}^{\mathrm{p}}\big)$ whenever $\ell^{\mathrm{r}} \geq \hat{\ell}^{\mathrm{r}}$.
5. *Standard part:* $\mathrm{st}(\ell) \equiv \mathbb{I}_{\ell^{\mathrm{r}}=0}\,\ell^{\mathrm{p}} \in \mathbb{R}_+$.

We identify each ℓ-number $(0, \ell^{\mathrm{p}}) \in L$ with the real number $\ell^{\mathrm{p}} \in \mathbb{R}_{++}$, so the previous operations are consistent with the usual addition, multiplication, and division of real numbers. Also, an ℓ-number $\ell \in L$ with $\mathrm{st}(\ell) = 0$ (i.e., $\ell^{\mathrm{r}} > 0$) is called an *infinitesimal.*

To gain some intuition about the operations between ℓ-numbers, fix two ℓ-numbers $\ell, \hat{\ell} \in L$ and let $(x_n)_n \in \ell$ and $(\hat{x}_n)_n \in \hat{\ell}$ be two sequences. The lexicographic order indicates which sequence tend to zero fastest. For example, $\ell > \hat{\ell}$ implies that either $\ell^{\mathrm{r}} < \hat{\ell}^{\mathrm{r}}$, hence $(x_n)_n$ tends to zero at a lower rate than $(\hat{x}_n)_n$, or that $\ell^{\mathrm{r}} = \hat{\ell}^{\mathrm{r}}$ and $\ell^{\mathrm{p}} > \hat{\ell}^{\mathrm{p}}$, so $(x_n)_n$ is asymptotically $\ell^{\mathrm{p}}/\hat{\ell}^{\mathrm{p}} > 1$ times $(\hat{x}_n)_n$. This implies:

$$\ell > \hat{\ell} \quad \Leftrightarrow \quad \lim_{n\to\infty} \frac{x_n}{\hat{x}_n} > 1 \ .$$

The sum of ℓ and $\hat{\ell}$ is such that $(x_n + \hat{x}_n)_n \in \ell + \hat{\ell}$. Intuitively, our construction only preserves the most likely terms: if $\ell^{\mathrm{r}} < \hat{\ell}^{\mathrm{r}}$, so $(x_n)_n$ tends to 0 at a slower rate than $(\hat{x}_n)_n$, then $(x_n + \hat{x}_n)_n$ tends to 0 at the same rate as $(x_n)_n$, hence $\ell + \hat{\ell} = \ell$. Multiplication and division are also simple and intuitive. We then have that, for all operations $\star \in \{+, \cdot, /\}$,

$$(x_n \star \hat{x}_n)_n \in \ell \star \hat{\ell}\ , \quad \text{that is,} \quad \lim_{n\to\infty} \frac{x_n \star \hat{x}_n}{n^{-(\ell\star\hat{\ell})^{\mathrm{r}}}\,(\ell\star\hat{\ell})^{\mathrm{p}}} = 1 \ .$$

The standard part of an ℓ-number ℓ is its closest real number: $\lim_{n\to\infty} x_n = \mathrm{st}(\ell)$.

3 ℓ-Strategy profiles and ℓ-equilibria

3.1 ℓ-Strategy profiles and assessments

In this section, we define two important concepts in this paper: ℓ-strategy profiles and the assessments they generate. Here, and in the rest of the paper, an extensive form game G is fixed.

Definition 3.1. An *ℓ-strategy profile* is a map $\lambda\colon A\to L$, such that, for each $I\in\mathcal{I}$:

1. $\sum_{a\in A^I}\lambda(a)=1$.
2. If $\iota(I)=0$ then $\lambda(a)=(0,\rho(a))$ for all $a\in A^I$.

The set of ℓ-strategy profiles is denoted Λ.

We can interpret an ℓ-strategy profile λ as the limit of a sequence of strategy $(\sigma_n)_n$. For example, we can define the sequence $(\sigma_n)_n$ as follows: for each $n\in\mathbb{N}$ and $a\in A$,

$$\sigma_n(a)\equiv\begin{cases} n^{-\lambda^{\mathrm{r}}(a)}\,\lambda^{\mathrm{p}}(a) & \text{if } \lambda^{\mathrm{r}}(a)>0,\\ M_n(\lambda,I^a)\,\lambda^{\mathrm{p}}(a) & \text{if } \lambda^{\mathrm{r}}(a)=0,\end{cases} \tag{3.1}$$

where $M_n(\lambda,I^a)$ is a factor that ensures that $\sum_{a'\in I^a}\sigma_n(a')=1$, and where the subindex n initialized so that $M_n(\lambda,I^a)\geq 0$ for all n and a (note that $\lim_{n\to\infty}M_n(\lambda,I^a)=1$). It is then clear that, for all $a\in A$, $(\sigma_n(a))_n\in\lambda(a)$.

Given an ℓ-strategy profile, we can easily define the likelihood of a history or an information set as follows. For each history $h\equiv(a_j)_{j=1}^J\in H$ and information set $I\in\mathcal{I}$, we define (with some abuse of notation)

$$\lambda(h)\equiv\textstyle\prod_{j=1}^J\lambda(a_j)\quad\text{and}\quad\lambda(I)\equiv\sum_{h'\in I}\lambda(h')\ .$$

The likelihood of a history is calculated in the same way as its probability: multiplying the likelihoods of the actions that compose it. Similarly, the likelihood of an information set is the sum of the likelihoods of the histories it contains.[6] Using $\sigma_n(h)$ and $\sigma_n(I)$ to denote the probabilities of h and I under the strategy profile in (3.1), it is easy to see that $(\sigma_n(h))_n\in\lambda(h)$ and $(\sigma_n(I))_n\in\lambda(I)$.

Assessments

We will now relate the concept of ℓ-strategy profile to the standard concept of assessment (Kreps and Wilson, 1982). Recall that an *assessment* is a pair (σ,μ), where $\sigma\in\Sigma$ is a

[6]Note that the likelihood of an information set is not affected by the likelihoods of histories which are infinitely less likely than other histories of the information set. If, for example, $I=\{h,h'\}$ and $\lambda^{\mathrm{r}}(h)<\lambda^{\mathrm{r}}(h')$ (so h is infinitely more likely than h'), then $\lambda(I)=\lambda(h)+\lambda(h')=\lambda(h)$.

strategy profile, and where $\mu : H\backslash T \to [0,1]$ is a belief system such that $\sum_{h\in I} \mu(h) = 1$ for all $I \in \mathcal{I}$. We can associate an assessment to an ℓ-strategy profile as follows:

Definition 3.2. The assessment *generated* by $\lambda \in \Lambda$ is the assessment $(\sigma^\lambda, \mu^\lambda)$ defined by $\sigma^\lambda(a) \equiv \mathrm{st}(\lambda(a))$ for all $a \in A$, and $\mu^\lambda(h) \equiv \mathrm{st}(\lambda(h)/\lambda(I))$ for all $I \in \mathcal{I}$ and $h \in I$.

Kreps and Wilson (1982) point out that not all assessments are plausible. Interpreting $\mu(h)$ as the belief that player $\iota(I)$ holds about history $h \in I$ at I, it is natural to require Bayes consistency if I is on path, and also to require that beliefs are updated consistently off path. An assessment (σ, μ) is *consistent* if there is a fully-mixed sequence $(\sigma_n)_n$ supporting (σ, μ); that is, $\sigma_n \to \sigma$ and $\sigma_n(h)/\sigma_n(I) \to \mu(h)$ for all $I \in \mathcal{I}$ and $h \in I$.

The following result establishes that consistent assessments coincide with assessments generated by ℓ-strategy profiles. Hence, it provides an easy way to generate consistent assessments without the need of computing limits of strategy profiles.

Proposition 3.1. *An assessment is consistent if and only if it is generated by some ℓ-strategy profile.*

Proposition 3.1 establishes that consistent assessments are generated by some ℓ-strategy profile, and that assessments generated by an ℓ-strategy profile are consistent. The "if" part of the statement follows from the fact that $(\sigma_n)_n$ defined in (3.1) supports $(\sigma^\lambda, \mu^\lambda)$, so $(\sigma^\lambda, \mu^\lambda)$ is a consistent assessment.

Proving the "only if" part of Proposition 3.1 is more involved. To see why, we fix some consistent assessment (σ, μ) and some sequence $(\sigma_n)_n$ supporting (σ, μ). Note that the rate of convergence of $\sigma_n(a)$ to $\sigma(a)$ may be very different for different actions $a \in A$, and not necessarily through a simple sequence of the form (3.1) for some $\lambda \in \Lambda$. For example, it could be that

$$\sigma_n(a) = 1/\log(n+1), \;\; \sigma_n(a') = e^{-n} + e^{-n^3}, \;\text{ and }\; \sigma_n(a'') = 3/\log(n+1)\, e^{-n} + e^{-n^2} \qquad (3.2)$$

for three different actions $a, a', a'' \in A$. Then, to find an ℓ-strategy profile generating (σ, μ), we need to assign an ℓ-number to each action in a way that the implied ℓ-numbers associated to histories represent their relative likelihood according to σ_n as $n \to \infty$. For example, if $h \equiv (a, a')$ and $h' \equiv (a'')$ happen to be the only histories in some information set $I \equiv \{h, h'\}$, then

$$\mu(h) = \lim_{n\to\infty} \frac{\sigma_n(a)\,\sigma_n(a')}{\sigma_n(a)\,\sigma_n(a') + \sigma_n(a'')} = \frac{1}{1+3} \; .$$

It then must be that

$$\lambda^{\mathrm{r}}(a) + \lambda^{\mathrm{r}}(a') = \lambda^{\mathrm{r}}(a'') \;\text{ and }\; \lambda^{\mathrm{p}}(a)\,\lambda^{\mathrm{p}}(a') = 3^{-1}\,\lambda^{\mathrm{p}}(a'') \; . \qquad (3.3)$$

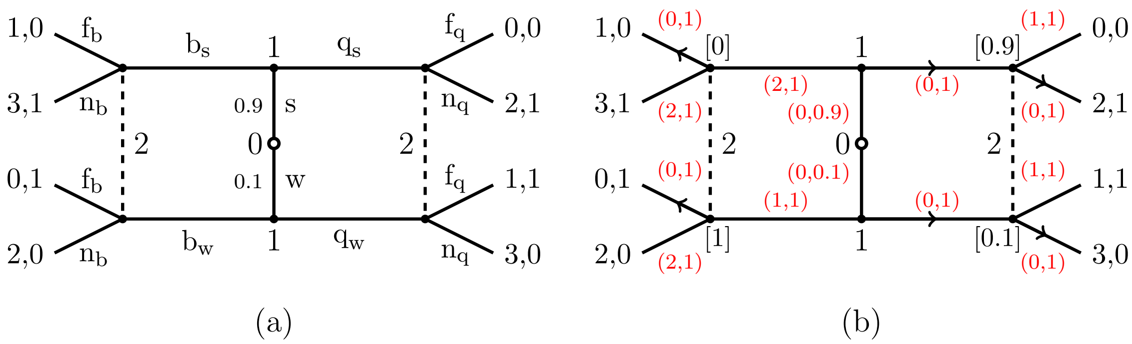

Figure 1: (a) Quiche-beer game and (b) an ℓ-strategy profile (in red). The arrows indicate the actions played under the ℓ-strategy profile, and the numbers in brackets are the beliefs of player 2 under the assessment generated by the ℓ-strategy profile.

The main difficulty of the construction is that many equations like (3.3) have to hold in a big extensive form game. Additionally, some inequality conditions must be satisfied, like for example $\lambda^{\mathrm{r}}(a)<\lambda^{\mathrm{r}}(a')$ since $\lim_{n\to\infty}(\sigma_n(a')/\sigma_n(a))=0$.

The proof of Proposition 3.1 provides an algorithm that provides, for each sequence $(\sigma_n)_n$ supporting some assessment, an ℓ-strategy profile that generates the same assessment. The algorithm overcomes the fact that ℓ-numbers cannot be independently assigned to actions. For example, a history $h''\equiv(a_j)_{j=1}^J$ with $\sigma_n(a_j)=1/\log(1+n)$ for all j is asymptotically infinitely more likely than history h' defined before, independently of how large J is. We use the fact that, given that the game is finite, it is enough to assign a likelihood rate to each a_j which is "small enough" so that $\sum_{j=1}^J \lambda^{\mathrm{r}}(a_j)<\lambda^{\mathrm{r}}(h'')$. Intuitively, the belief system μ^λ is only determined by the likelihoods of the histories which are infinitely more likely than the rest of the histories in each information set, and not the exact speeds at which the likelihoods converge.

Remark 3.1. A corollary of Proposition 3.1 is that the *set of simple strategy profiles* (i.e., of the form (3.1)) is rich enough to generate *all* consistent assessments. This simplifies proving or disproving the consistency of an assessment: if an assessment is consistent, it is supported by some simple strategy profile. There are, however, two main advantages of using ℓ-strategy profiles. The first is that, unlike simple sequences, ℓ-numbers are closed to addition. Hence, the likelihood of an information set I, which is the addition of the likelihoods of its histories, is an ℓ-number as well, while $\sigma_n(I)$ is typically a sum of various terms. The second advantage is that the requirement that the likelihoods assigned to the actions available at a given information set (condition 1 in Definition 3.1) imposes no restriction on the actions played with infinitesimal likelihood. Hence, terms such as $M_n(\lambda, I^a)$ in equation (3.1), which complicate operating with simple sequences of strategy profiles, play no role when using ℓ-strategy profiles.

Example 3.1. Figure 1 depicts the beer-quiche game and an ℓ-strategy profile.

3.2 ℓ-Equilibria and sequential equilibria

In this section, we introduce the concept of ℓ-equilibrium and its relationship to the concept of sequential equilibrium.

Payoffs and ℓ-equilibria

Before defining optimality, we derive the payoff a player receives from playing a given action at an information set. Fix an ℓ-strategy profile λ. For each $a \in A$, we let

$$u(a|\lambda) \equiv \sum_{t \in T^a} \mathrm{st}\left(\frac{\lambda(t)}{\lambda(I^a)\, \lambda(a)} \right) u_{\iota(I^a)}(t) \tag{3.4}$$

be player $\iota(I^a)$'s payoff conditional on I^a being reached and a being played, where $T^a \subset T$ is the set of terminal histories that contain a as one of its elements (we omit the subindex $\iota(I^a)$ to ease notation). The term $\lambda(t)/(\lambda(I^a)\,\lambda(a))$ in the previous expression is the likelihood of t conditional on I^a being reached and player $\iota(I^a)$ choosing a. Indeed, for every history $h \in I^a$, we have that $\lambda(h) \leq \lambda(I^a)$. This implies that, since any $t \in T^a$ can be written as $(h, a, a_1, ..., a_J)$ for some $h \in I^a$ and some $a_1, ..., a_J \in A$, we have that

$$\lambda(t) = \lambda(h)\, \lambda(a)\, \textstyle\prod_{j=1}^{J} \lambda(a_j) \leq \lambda(I^a)\, \lambda(a)\ .$$

In fact, it is easy to see that $\sum_{t \in T^a} \lambda(t) = \lambda(I^a)\, \lambda(a)$, which implies that

$$\sum_{t \in T^a} \frac{\lambda(t)}{\lambda(I^a)\, \lambda(a)} = 1\ .$$

We proceed by defining the concept of ℓ-equilibrium:

Definition 3.3. λ is an *ℓ-equilibrium* if $\mathrm{st}(\lambda(a)) > 0$ implies $u(a|\lambda) \geq u(a'|\lambda)$ for all $a \in A$ and $a' \in I^a$. We denote the set of ℓ-equilibria as Λ^*.

An ℓ-equilibrium is an ℓ-strategy profile where the only actions which are played with positive probability are sequentially optimal, that is, actions that maximize $u(\cdot|\lambda)$ among all other actions available at the information set they are available in.

Relationship to sequential equilibria

We now relate the concept of ℓ-equilibrium to the well-known concept of sequential equilibrium, first defined in Kreps and Wilson (1982). They define an assessment (σ, μ) to be a *sequential equilibrium* if it is consistent (as defined before) and *sequentially rational*, that is, satisfying that for each $a \in A$:

$$\sigma(a) > 0 \quad \Rightarrow \quad u(a|\mu, \sigma) \geq u(a'|\mu, \sigma) \ \text{ for all } a' \in A^{I^a}\ ,$$

where

$$u(a|\mu,\sigma) \equiv \sum_{t\in T^a} \Pr^{\mu,\sigma}(t|a)\, u_{\iota(I^a)}(t) \ ,$$

and where for each $t \in T^a$ of the form $t=(h,a,a_1,...,a_J)$ we have

$$\Pr^{\mu,\sigma}(t|a) \equiv \mu(h) \prod_{j=1}^{J} \sigma(a_j) \ .$$

The following result illustrates the close connection between sequential equilibria and ℓ-equilibria.

Proposition 3.2. (σ,μ) *is a sequential equilibrium if and only if it is generated by some ℓ-equilibrium. Hence, an ℓ-equilibrium exists.*

Proposition 3.2 exemplifies the usefulness of our approach. Indeed, proving that a given assessment is a sequential equilibrium is not easy in many applications. The difficulty typically resides in finding a sequence of fully-mixed assessments that converges to (σ,μ). Such difficulties have favored the use of equilibrium concepts which are not as powerful in selecting Nash equilibria, such as perfect Bayesian equilibria (Fudenberg and Tirole, 1991b). ℓ-Equilibria provide a simple characterization of sequential equilibria: any ℓ-strategy profile satisfying sequential rationality is a sequential equilibrium.

The first statement of Proposition 3.2 follows from Proposition 3.1, which guarantees the consistency of generated assessments, and the fact that $u(a|\lambda)=u(a|\mu^\lambda,\sigma^\lambda)$, which makes the second condition in Definition 3.3 equivalent to sequential rationality. The proof of existence is then trivial: since a sequential equilibrium exists (by Kreps and Wilson, 1982), and since it is generated by some ℓ-strategy profile (by Proposition 3.1), such ℓ-strategy profile is also an ℓ-equilibrium.

Example 3.2 (Continuation of Example 3.1). It is easy to see that the ℓ-strategy profile depicted in Figure 1 (b) is an ℓ-equilibrium, hence generates a sequential equilibrium.

4 ℓ-Trembles and their associated ℓ-equilibria

Our next goal is to study stability using ℓ-numbers. As stability is based on perturbations of the game, we introduce the concepts of ℓ-tremble and ℓ-equilibrium for an ℓ-tremble in this section. We will show that these concepts are related to the standard concepts of tremble, trembling-hand equilibria and ε-perfection. Section 5 will use these concepts and results to define and analyze ℓ-stability.

4.1 ℓ-Trembles and trembles

Recall that a *(behavioral) tremble* is a sequence $(\eta_n : A \to (0,1])_n$ satisfying that that (i) $\lim_{n\to\infty} \eta_n(a)=0$ for all $a\in A$, and (ii) $\sum_{a\in A^I} \eta_n(a) \leq 1$ for all $I\in\mathcal{I}$ and $n\in\mathbb{N}$. Differently from a strategy profile, each η_n is not required to be a probability distribution when restricted to a given information set I, and instead we require $\eta_n(a)$ to be strictly positive for all a and to converge to 0 as $n \to \infty$ (and also that $\sum_{a \in A^I} \eta_n(a) \leq 1$). The concept of ℓ-tremble is then analogous to the concept of tremble.

Definition 4.1. An *ℓ-tremble* is a map $\tilde{\lambda}: A \to L$ such that $\text{st}(\tilde{\lambda}(a)) = 0$ for all $a \in A$. The set of ℓ-trembles is $\tilde{\Lambda}$.

In the same way that trembles are not sequences of strategy profiles, the concept of ℓ-tremble relaxes the conditions that an ℓ-strategy profile satisfies: an ℓ-tremble differs from an ℓ-strategy profile in that each action is asinged an infinitesimal, and hence ℓ-trembles are not probability distributions when restricted to the actions available at a given information set. Trembles and ℓ-trembles are interpreted as representations of small likelihoods with which players make mistakes.

4.2 ℓ-Equilibria for an ℓ-tremble

We now define and characterize ℓ-equilibria for a given ℓ-tremble. This equilibrium concept will be the basis of our definition of ℓ-stable ourcome.

Definition 4.2. $\lambda \in \Lambda$ is an *ℓ-equilibrium for* $\tilde{\lambda} \in \tilde{\Lambda}$ if, for all $I \in \mathcal{I}$ and $a \in A^I$,

1. **Superseding:** $\lambda(a) \geq \tilde{\lambda}(a)$.
2. **Optimality:** $\lambda(a) > \tilde{\lambda}(a)$ only if $u(a|\lambda) \geq u(a'|\lambda)$ for all $a' \in A^I$.

We denote the set of ℓ-equilibria for $\tilde{\lambda}$ as $\Lambda^*(\tilde{\lambda})$.

The concept of ℓ-equilibria for a given ℓ-tremble shares the spirit of trembling-hand perfect equilibria (Selten, 1975), but focusses on a particular tremble. Our requirement of "superseding" hints at the interpretation of the ℓ-tremble as an asymptotic probability of making mistakes: conditional on a given information set being reached, the likelihood that the corresponding player can assign to an action a can not be lower than $\tilde{\lambda}(a)$. Unlike the concept of ℓ-equilibrium, we now require sequential optimality for some actions with infinitesimal likelihood: the ones with a strictly higher likelihood under the ℓ-equilibrium than under the ℓ-tremble.[7] As a result, the set of ℓ-equilibria for a given perturbation is typically smaller than the set of ℓ-equilibria (see Examples 6.1 and 6.2 below). The following result establishes that, nonetheless, the set of ℓ-equilibria for some ℓ-tremble coincides with the set of ℓ-equilibria.

[7] Apart from its use in studying ℓ-stability, the concept of ℓ-equilibrium for a given ℓ-tremble is useful when there are some particular trembles of economic interest, for example for economic reasons (e.g., people with lower financial literacy trembling more frequently in financial decisions).

Proposition 4.1. *An ℓ-strategy profile is an ℓ-equilibrium if and only if it is an ℓ-equilibrium for some ℓ-tremble; that is,* $\Lambda^* = \cup_{\tilde{\lambda}\in\tilde{\Lambda}}\Lambda^*(\tilde{\lambda})$.

The "if" part of Proposition 4.1 follows from the observation that each ℓ-equilibrium λ is an ℓ-equilibrium for the ℓ-tremble $\tilde{\lambda}$ defined as $\tilde{\lambda}(a) \equiv \lambda(a)$ if $\mathrm{st}(\lambda(a)) = 0$ and $\tilde{\lambda}(a) \equiv (1,1)$ otherwise. The "only if" part is implied by the fact that, if λ is an ℓ-equilibrium for some $\tilde{\lambda}$, then $\mathrm{st}(\lambda(a)) > 0$ only if a is sequentially optimal, but then this implies λ is also an ℓ-equilibrium.

Relationship to trembling-hand and ε-perfection

In this subsection, we relate the concept of ℓ-equilibrium to the concepts of trembling-hand and perfect ε-equilibrium. Doing so will be useful to relate the concepts of ℓ-stability and stability in Section 5.

Definition 4.3. $\sigma_n \in \Sigma$ is an *ε-perfect equilibrium for η_n* if, for all $I \in \mathcal{I}$ and $a \in A^I$,

1. $\sigma_n(a) \geq \eta_n(a)$.
2. $\sigma_n(a) > \eta_n(a)$ only if $u(a|\sigma_n) \geq u(a'|\sigma_n) - \varepsilon$ for all $a' \in A^I$.

We use $\Sigma^*_\varepsilon(\eta_n)$ denote the set of ε-perfect equilibria for η_n.

The definition of an ε-perfect equilibrium for a given tremble is similar to the definition of ℓ-equilibrium for an ℓ-tremble.[8] Both concepts require "superseding" (actions should be played with at least the likelihood specified by the tremble/ℓ-strategy profile) and "optimality" (actions which are played with a higher likelihood than the tremble cannot be (ε-)suboptimal). As we will see, the following concept is analogous to the concept of ℓ-equilibrium for some ℓ-tremble:

Definition 4.4. (σ, μ) is a *perfect equilibrium for* $(\eta_n)_n$ if there is a sequence $(\sigma_n)_n$ supporting (σ, μ), and a sequence $(\varepsilon_n)_n \to 0$ such that $\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta_n)$ for all n.

The following proposition provides an equivalence result between the concepts of ℓ-equilibrium for some ℓ-tremble and perfect equilibrium for some tremble.

Proposition 4.2. *Let λ be an ℓ-strategy profile. Then, λ is an ℓ-equilibrium for some ℓ-tremble if and only if $(\sigma^\lambda, \mu^\lambda)$ is a perfect equilibrium for some tremble.*

[8] Note that we require ε-optimality for the strategy at each information set. Hence, our equilibrium concept is analogous to contemporaneous perfect ε-equilibria defined in Mailath, Postlewaite, and Samuelson (2005).

Trembles associated to ℓ-trembles

While Proposition 4.2 establishes that any ℓ-equilibrium for some ℓ-tremble generates a perfect equilibrium for some tremble, it does not specify how the ℓ-tremble and the tremble are related to each other. To relate these objects, it is convenient to first associate a tremble to each ℓ-tremble.

Definition 4.5. The tremble *associated* to $\tilde{\lambda}\in\tilde{\Lambda}$, denoted $(\eta_n^{\tilde{\lambda}})_n$, is defined as $\eta_n^{\tilde{\lambda}}(a)\equiv n^{-\tilde{\lambda}^{\mathrm{r}}(a)}\,\tilde{\lambda}^{\mathrm{p}}(a)$ for all $a\in A$ and $n\in\mathbb{N}$.[9]

The following results establishes a close relationship between perfect equilibria for a tremble associated to some ℓ-tremble and ℓ-equilibria for such ℓ-tremble.

Proposition 4.3. (σ,μ) *is a perfect equilibrium for* $(\eta_n^{\tilde{\lambda}})_n$ *if and only if it is generated by some* ℓ*-equilibrium for* $\tilde{\lambda}$. *Hence, for each* $\tilde{\lambda}\in\tilde{\Lambda}$, *an* ℓ*-equilibrium for* $\tilde{\lambda}$ *exists.*

The "if" part of Proposition 4.3 is proven by assuming that there is an ℓ-equilibrium λ for $\tilde{\lambda}$. It then follows that, for all $\varepsilon>0$, σ_n defined in equation (3.1) is an ε-equilibrium for $\eta_n^{\tilde{\lambda}}$ if n is large enough. Note that, in general, $\sigma_n\not\in\Sigma_0^*(\eta_n^{\tilde{\lambda}})$; that is, σ_n is not a Nash equilibrium of a game where players tremble according to $\eta_n^{\tilde{\lambda}}$. In fact, it may be that $u(a|\sigma_n)<u(a'|\sigma_n)$ for all n but $\lim_{n\to\infty}(u(a|\sigma_n)-u(a'|\sigma_n))=0$ for some a and a' in the same information set, and yet that $\lambda(a)>0$, as $u(a|\lambda)=u(a'|\lambda)$ makes a sequentially rational in the limit.[10]

Proving the "only if" part in Proposition 4.3 is more involved. Let (σ,μ) be a perfect equilibrium for $(\eta_n^{\tilde{\lambda}})_n$, hence there exists some sequence $(\sigma_n\in\Sigma_{1/n}^*(\eta_n^*))_n$ supporting (σ,μ), but such sequence is not necessarily simple (of the form (3.1)). Hence, to find some ℓ-equilibrium λ for $\tilde{\lambda}$, we proceed similarly as in the proof of Proposition 3.1, with the additional constraint that $\lambda(a)=\tilde{\lambda}(a)$ whenever a is suboptimal. We show that the methodology developed in the proof of Proposition 3.1 can be adapted to obtain a ℓ-equilibrium for $\tilde{\lambda}$. Existence of an ℓ-equilibrium for $\tilde{\lambda}$ then follows from the fact that, for each n, there is some $\sigma_n\in\Sigma_0^*(\eta_n^{\tilde{\lambda}})$ (which is a Nash equilibrium of the perturbed game). Then, taking a subsequence if necessary, $(\sigma_n)_n$ supports some assessment (σ,μ), which is a perfect equilibrium for $(\eta_n^{\tilde{\lambda}})_n$.

---

[9]We let the sequence $(\eta_n^{\tilde{\lambda}})_n$ begin at an index large enough that $\sum_{a\in A^I}\eta_n(a)\leq 1$ for all n.

[10]Kreps and Wilson (1982) show that similar considerations apply to the study of sequential equilibria. In their words, "we require optimality only 'at the limit,' while Selten requires optimality approaching the limit, is what is significant in terms of tractability and mathematical properties." In fact, it is not difficult to see that σ is part of a sequential equilibrium if and only if there exist sequences $(\varepsilon_n)_n$ converging to 0 and $(\sigma_n\in\Sigma_{\varepsilon_n}^*(0))_n$ completely mixed and converging to σ.

5 ℓ-Stable outcomes

In this section, we use the definitions and results in Sections 3 and 4 to introduce the concept of ℓ-stable outcomes, prove their existence and provide characterization results.

5.1 Stable outcomes

We begin recalling the definition of outcome. An *outcome* o (of G) is a probability distribution over terminal histories, so $o \in \Delta(T)$. Each strategy profile σ *generates* a unique outcome o^σ, where each $(a_j)_{j=1}^J \in T$ is assigned probability $o^\sigma((a_j)_{j=1}^J) = \prod_{j=1}^J \sigma(a_j)$. Similarly, an ℓ-strategy profile λ *generates* a unique outcome o^λ, where the probability of each $t \in T$ is $o^\lambda(t) \equiv \mathrm{st}(\lambda(t))$. Note that $o^\lambda = o^{\sigma^\lambda}$.

There are different reasons that make the use of equilibrium outcomes appealing (compared to using equilibria). First, outcomes fully describe on-path behavior, hence they constitute the basis for the predictions that a game provides. Even though off-path behavior is important to determine the optimality of on-path behavior, different equilibria with the same outcome provide the same predictions. Second, they reduce spurious multiplicity arising from off-path behavior. Last, but not least, while equilibria are generally fragile to small perturbations of the game, there are equilibrium outcomes which are robust to both trembles and payoff perturbations.

The study of stable outcomes begins with Kohlberg and Mertens (1986). They show that the set of Nash equilibria of a game has a finite number of connected components and, generically in the payoffs, all equilibria in a connected component have the same outcome.[11] Furthermore, also generically in the payoffs, at least one of these outcomes o is *stable*: if the game is perturbed (through a small tremble), the resulting game has an equilibrium outcome which is close to o. In our notation, their result establishes that, for generic u, there is an outcome o such that, for any tremble $(\eta_n)_n$, there is a sequence $(\sigma_n \in \Sigma_0^*(\eta_n))_n$ such that $(o^{\sigma_n})_n$ converges to o. Kohlberg and Mertens show that stable outcomes satisfy backwards induction, iterated dominance, and invariance.

There are two important drawbacks that make using stable outcomes difficult in practice: existence and tractability. Even in the ideal case where an economist correctly guesses that a given outcome o is the unique stable outcome of a game of interest, proving so is difficult in general. S/he has to prove two facts. First, s/he has to prove "uniqueness" by showing that o is the unique candidate to be a stable outcome. This can be proven if, for example, s/he (1) is able to find a convenient tremble such that the set of equilibria of each of the corresponding perturbed games can be computed, and

[11] We use the usual distance between strategy profiles (sup-norm on $\mathbb{R}^A$), the usual distance between utility functions (sup-norm on $\mathbb{R}^{N\times T}$), and the usual distance between outcomes (sup-norm on $\mathbb{R}^T$).

(2) show that all sequences of equilibrium outcomes of the perturbed games along the sequence converge to o.[12] Then, the economist has to prove "existence" by showing that any other tremble has an associated sequence of equilibria with outcomes converging to o.[13] This is, in general, much more difficult to prove, as studying the sets of equilibria of the perturbed game for an arbitrary tremble is often not feasible. (An alternative would be characterizing stable sets of equilibria, which exist in all games, but this is typically equally infeasible.)

5.2 ℓ-Stable outcomes

Here, we provide the concept of ℓ-stable outcome, show its existence for any game, and characterize its relationship to stable outcomes. In Section 5.3, we provide some properties of ℓ-stable outcomes and illustrate how they can be use to prove a given outcome is stable.

Definition 5.1. An outcome o is *ℓ-stable* if, for each ℓ-tremble $\tilde{\lambda}$, there is an ℓ-equilibrium for $\tilde{\lambda}$ generating o.

The concept of ℓ-stable outcome is analogous to that of stable outcome, now using ℓ-trembles instead of trembles. The following result establishes that, differently from stable outcomes, ℓ-stable outcomes exist in all games, with no genericity requirement on the payoffs. We keep the proof in the main text as it is important and concise.

Proposition 5.1. *There is an ℓ-stable outcome.*

Proof. For a given utility function $\hat{u} : N \to \mathbb{R}$, we let $G(\hat{u})$ be the game defined in Section 2.1 with utility function given by $\hat{u}$ instead of u. Let $(\hat{u}_n)_n$ be a sequence of utility functions converging to u such that, for each n, $G(\hat{u}_n)$ has a stable outcome denoted o_n. Note that, since a stable outcome exists for generic utility functions (by Kohlberg and Mertens, 1986), a sequence $(\hat{u}_n)_n$ with the previous properties exists. Taking a subsequence if necessary, assume that $(o_n)_n$ converges to some outcome o. We will prove that o is ℓ-stable.

Fix some sequence $(\varepsilon_n)_n \searrow 0$ and $\tilde{\lambda} \in \tilde{\Lambda}$. For each $n' \in \mathbb{N}$, we let $\big(\sigma^{\tilde{\lambda}}_{n,n'} \in \Sigma^*_0(\eta^{\tilde{\lambda}}_n, \hat{u}_{n'})\big)_n$ be a sequence of Nash equilibria with outcomes converging to o_n, which exists by the stability of o_n. We now argue that, for each n, there is some j_n such that $\sigma^{\tilde{\lambda}}_{n,j_n}$ is an

[12]More generally, for each candidate o' of being a stable outcome (e.g., outcomes of sequential equilibria), the economist should find a tremble with no outcome sequence converging to o'.

[13]Many games of interest do not feature generic payoffs. For example, in most reputation and bargaining games, the player's payoff equals to the discounted payoff of a stage game (at the moment of agreement for bargaining games), using some discount factor (constant over time, and typically equal across players). Hence, stable outcomes are typically not guaranteed to exist.

ε_n-equilibrium for $\eta_n^{\tilde{\lambda}}$ (in the game with payoff function u). Superseding is satisfied for all n, since $\sigma_{n,j_n}^{\tilde{\lambda}} \in \Sigma_0^*(\eta_n^{\tilde{\lambda}}, \hat{u}_{j_n})$. Also, for each $a \in A$ with $\sigma_{n,j_n}^{\tilde{\lambda}}(a) > \eta_n(a)$, we have

$$\hat{u}_n(a|\sigma_{n,j_n}^{\tilde{\lambda}}) \geq \hat{u}_n(a'|\sigma_{n,j_n}^{\tilde{\lambda}}) \text{ for all } a' \in I^a.$$

Hence, if j_n is chosen so that $\|u - \hat{u}_{j_n}\| < \varepsilon_n/2$, we have

$$u(a|\sigma_{n,j_n}^{\tilde{\lambda}}) \geq u(a'|\sigma_{n,j_n}^{\tilde{\lambda}}) - \varepsilon_n \ .$$

Taking a subsequence if necessary, assume that $(\sigma_{n,j_n}^{\tilde{\lambda}})_n$ supports some assessment (σ, μ). Clearly, (σ, μ) is a perfect equilibrium for $(\eta_n)_n$ with outcome o. By Proposition 4.3, (σ, μ) is generated by some ℓ-equilibrium for $\tilde{\lambda}$. □

Proposition 5.1 is important as it shows that ℓ-stability overcomes an important drawback of working with stable outcomes in applications: having to prove that they exist. As we will see, our existence result simplifies finding ℓ-stable outcomes, as they are often obtained by ruling out other candidates.

Relationship to stable outcomes

Using Proposition 4.3, we can provide the following characterization of ℓ-stable outcomes in terms of perfect equilibrium outcomes.[14]

Proposition 5.2. *An outcome o is ℓ-stable if and only if, for any $\tilde{\lambda} \in \tilde{\Lambda}$, o is a perfect equilibrium outcome for $(\eta_n^{\tilde{\lambda}})_n$.*

A first implication of Proposition 5.2 is that the concept of ℓ-stable outcome is not powerful in selecting equilibria of normal-form games. Such observation follows from the analysis in Jackson, Rodriguez-Barraquer, and Tan (2012). They first define a strategy profile σ in a normal-form game to be a *trembling*-hand perfect equilibrium* if, for any tremble $(\eta_n)_n$, there exist two sequences $(\varepsilon_n)_n \to 0$ and $(\sigma_n \in \Sigma_{\varepsilon_n}^*(\eta_n)) \to \sigma$. They subsequently show that all Nash equilibria of any finite normal form game are trembling*-hand perfect equilibria. An immediate implication is that all Nash outcomes of a normal form game are ℓ-stable. We will see that, the set of ℓ-stable outcomes in extensive form games with incomplete information may be significantly smaller than the set of Nash equilibrium outcomes.[15]

[14]Analogously to Definition 4.4, we say that o is a *perfect equilibrium outcome* for $(\eta_n)_n$ if there is a sequence $(\sigma_n)_n$ and a sequence $(\varepsilon_n)_n \searrow 0$ such that $\sigma_n \in \Sigma_{\varepsilon_n}^*(\eta_n)$ for all n and $o^{\sigma_n} \to o$.

[15]Similarly, Fudenberg and Tirole (1991a) show that, in normal-form games, all and only the outcomes of Nash equilibria are robust to payoff perturbations. Nonetheless, Takahashi and Tercieux (2020) show that payoff robustness has a significant selection power in extensive form games.

A second implication of Proposition 5.2 is that ℓ-stability is a less restrictive concept than stability. Indeed, the crucial difference between ℓ-stable outcomes and stable outcomes is the use of ε-perfect equilibria versus the use of Nash equilibria in the perturbed games. The following corollary establishes that, as a result, the set of stable outcomes is smaller than the set of ℓ-stable outcomes.

Corollary 5.1. *If a stable outcome exists, it is ℓ-stable.*

It is clear that the converse of Corollary 5.1 is not true in general: while all games have ℓ-stable outcomes, some do not have any stable outcome. A sort of converse result is nevertheless true when there is a unique ℓ-stable outcome. Intuitively, if a game had a unique ℓ-stable outcome which was not stable, there would be a tremble $(\eta_n)_n$ such that, for n large enough, no equilibrium of the perturbed game would have an equilibrium close to o. The proof shows that, similarly as in the proof of Proposition 5.1, we could then construct an alternative ℓ-stable outcome, contradicting the assumption that o was the unique ℓ-stable outcome.

Proposition 5.3. *If a unique ℓ-stable outcome exists, it is the unique stable outcome.*

We end this subsection by pointing that ℓ-stable outcomes are robust not only to perturbations in the players' behavior, but also in the player's payoff. The following is a generalization of Proposition 5.2.

Proposition 5.4. *An outcome o is ℓ-stable if and only if, for any $\tilde{\lambda} \in \tilde{\Lambda}$ and $(u_n)_n$ converging to u, there are two sequences $(\varepsilon_n)_n$ and $(\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta^{\tilde{\lambda}}_n, u_n))_n$ with $o^{\sigma_n} \to o$ (where $\Sigma^*_{\varepsilon_n}(\eta^{\tilde{\lambda}}_n, u_n)$ is the set of ε_n-equilibria for $\eta^{\tilde{\lambda}}_n$ in the game with payoffs $u_n : T \to \mathbb{R}^N$).*

5.3 Obtaining ℓ-stable and stable outcomes

In this section, we provide some properties of ℓ-stable outcomes. As the following theorem shows, an ℓ-stable outcome satisfies versions of the definitions of invariance, iterated dominance and forward induction in Kohlberg and Mertens (1986). We keep the proof of forward induction in the main text as it is simple and intuitive.

Proposition 5.5. *Let o be an ℓ-stable outcome.*[16]

1. ***Forward induction:*** *Assume $I \in \mathcal{I}$ is on path of o and $a \in A^I$ is such that*

$$\max_{\lambda \in \Lambda^*(o)} u(a|\lambda) < u(I|o) \ ,$$

[16]In the statement of this proposition, we abuse language by using "equivalent outcomes" without defining them. Nevertheless, their definition (which depends on each property) is straightforward.

where $u(I|o)$ is player $\iota(I)$'s payoff under o conditional on I being reached, and $\Lambda^*(o)$ is the set of ℓ-equilibria with outcome o. Then, the game where a is eliminated (and all histories following it) has an ℓ-stable outcome equivalent to o.[17]

2. ***Iterated strict dominance:*** *If a strictly dominated action is eliminated (and all histories following it), the resulting game has an ℓ-stable outcome equivalent to o.*

3. ***Invariance to reordering simultaneous moves:*** *If $I' = I \times A^I$ (i.e., I and I' are simultaneous), a game where the order of I and I' is reversed has an ℓ-stable outcome equivalent to o.*

Proof of "forward induction" (the other properties are proven in Appendix A). Let o be an ℓ-stable outcome. Let I be an on-path information set and $\hat{a} \in A^I$ provide a payoff strictly lower than the outcome's payoff of player $\iota(I)$ at I, hence $\hat{a}$ not played in o. Let G' denote the game where $\hat{a}$ (and all consecutive histories) is eliminated, and $A' \subset A \setminus \{\hat{a}\}$ be its set of actions. Let $\tilde{\lambda}'$ be an ℓ-tremble of G', and $\underline{\tilde{\lambda}}^{\mathrm{r}} \equiv \max\{\tilde{\lambda}'^{\mathrm{r}}(a') | a' \in A'\}$. Define $\tilde{\lambda}$ as follows:

$$\tilde{\lambda}(a) \equiv \begin{cases} \tilde{\lambda}'(a) & \text{if } a \in A', \\ (\underline{\tilde{\lambda}}^{\mathrm{r}} + 1, 1) & \text{otherwise,} \end{cases}$$

for all $a \in A$, and note that $\tilde{\lambda}$ is an ℓ-tremble of G. Let λ be an ℓ-equilibrium for $\tilde{\lambda}$ with outcome o (which exists since o is ℓ-stable). We claim that, when restricted to A', λ is an ℓ-equilibrium for $\tilde{\lambda}'$ on G' with outcome o. Note that, since $u(\hat{a}|\lambda) < u(I^{\hat{a}}|o)$, we have that $\lambda(\hat{a}) = (\underline{\tilde{\lambda}}^{\mathrm{r}} + 1, 1)$, which implies that each terminal history containing $\hat{a}$ is infinitely less likely than any of the histories not containing $\hat{a}$. It is then easy to see that the ℓ-strategy profile of G' defined by $\lambda' \equiv \lambda|_{A'}$ is an ℓ-equilibrium for $\tilde{\lambda}'$ (in G') with outcome o. □

It is not difficult to see that an ℓ-stable outcome may fail *admissibility*, that is, players may play weakly dominated actions on the path of play. This is not surprising, since admissibility is known to be incompatible with iterated (strict) dominance. Also, the "iterated strict dominance" property permits using *backwards induction* arguments. Indeed, it follows from Proposition 5.5 that an action which is strictly dominated, or which is strictly suboptimal in under any ℓ-equilibrium, can be eliminated without affecting the set of ℓ-stable outcomes. In particular, an extensive form game of complete information with generic payoffs has a unique ℓ-stable (and hence stable) outcome.

[17] The result also holds if $\max_{t \in T^a} u_{\iota(I)}(t) < u(I|o)$, a more restrictive but easier to verify condition.

Remark 5.1. It is easy to see that forward induction and iterated strict dominance can be generalized to the same property: an ℓ-stable outcome o remains ℓ-stable if an action which is *strictly dominated under* o is eliminated; that is, if $a \in A$ such that there is some $a' \in A^{I^a}$ with $u(a'|\lambda) > u(a|\lambda)$ for all $\lambda \in \Lambda^*(o)$ is eliminated. This more general property is closer to the definition of forward induction in Kohlberg and Mertens (1986, Proposition 6).

Obtaining ℓ-stable outcomes

Here, we provide some strategies an economist can use to obtain ℓ-stable outcomes.

In many games with asymmetric information, the fact that actions endogenously signal private information constitutes a main source of equilibrium multiplicity. These games tend to have a unique stable outcome if some sort of "single-crossing" condition holds, even when payoffs are not "generic". In simple signaling games á la Spence (1973), for example, such stable outcome is the least-costly, fully-separating outcome, called the *Riley outcome* (see Riley, 1979). The following corollary of Propositions 5.1 and 5.3 implies that, in such cases, the economist only needs to obtain one "right" ℓ-tremble to prove that o is the unique ℓ-stable outcome, and hence it is stable:

Corollary 5.2. *If there is an ℓ-tremble $\tilde{\lambda}$ such that there is a unique ℓ-equilibrium outcome for $\tilde{\lambda}$, then such outcome is the unique stable outcome.*

As Examples 6.1 and 6.2 below illustrate, an ℓ-tremble that helps proving uniqueness of an ℓ-stable outcome in signaling games satisfying single-crossing usually involves "high types" (the types with lowest signaling cost) trembling with a higher likelihood than low types.

Games not satisfying a single-crossing condition may have multiple ℓ-stable outcomes. Nevertheless, through one or more ℓ-trembles, the economist may able to reduce the set of outcome candidates to some set $O^\dagger$. If all outcomes in $O^\dagger$ satisfy a given property (full/partial separation, delay, etc), then the economist can claim that such property is "ℓ-stable", in the sense that all ℓ-stable outcomes satisfy it.

Example 5.1 (Continuation of Example 3.2). We argue that the outcome in Figure 1(b) is **not** ℓ-stable using forward induction and iterated strict dominance. If it were to be ℓ-stable, it would remain ℓ-stable when action $\mathrm{b_w}$ is eliminated, since the maximum payoff the weak type can achieve by playing it is lower than her payoff from playing $\mathrm{b_s}$ under the outcome. In the game without action $\mathrm{b_w}$, action $\mathrm{f_b}$ is strictly dominated, hence can be eliminated as well. In the resulting game, the strong type prefers playing $\mathrm{b_s}$ to playing $\mathrm{q_s}$, contradicting the ℓ-stability of the outcome. Then, the outcome in Figure 1(b) is not ℓ-stable. Since there is only one other outcome of a sequential equilibrium (where player 1 plays both $\mathrm{b_s}$ and $\mathrm{b_w}$ for sure, and player 2 plays $\mathrm{n_b}$ for sure), this other outcome is the unique ℓ-stable outcome (by Proposition 5.1), and hence it is the unique stable outcome as well (by Proposition 5.3).

6 Examples

In this section, we provide some examples of how to use our results. Due to obvious size constraints of the current paper, the examples are simple, and other approaches (e.g., brute force) could used to obtain stable outcomes. See Dilmé (2021) for an example of a game where brute force does not work.

Example 6.1 (Continuation of Example 5.1). To show how ℓ-trembles can be used to contradict the ℓ-stability of an outcome, we now argue that the outcome in Figure 1(b) is not ℓ-stable without using forward induction (as in Example 5.1). To see this, consider an ℓ-tremble $\tilde{\lambda}$ where $\tilde{\lambda}^{\mathrm{r}}(\mathrm{b_s}) < \tilde{\lambda}^{\mathrm{r}}(\mathrm{b_w})$ (that is, player 1 trembles to beer infinitely more often when she is strong). Let λ be an ℓ-equilibrium for $\tilde{\lambda}$. If $\mu^{\lambda}_{\mathrm{b}}(\mathrm{s},\mathrm{b_s}) > 1/2$, then player 2 does not fight with probability one after beer, but then player 1 has a profitable deviation if she is strong. Also, $\mu^{\lambda}_{\mathrm{b}}(\mathrm{s},\mathrm{b_s}) \leq 1/2$ only if player 1 finds it optimal to choose beer when she is weak, but this is clearly suboptimal given that her payoff from doing so is lower than the payoff from choosing quiche. Hence, there is no equilibrium for $\tilde{\lambda}$ with the outcome in Figure 1 (b), so it is **not** ℓ-stable.

Example 6.2 (Signaling). In this example we consider a version of the Spence (1973) model. Nature first decides the type of player 1, $\theta \in \{L, H\}$. Then player 1 chooses the effort $e \in \mathcal{E} \equiv \{0, \Delta, 2\Delta, ..., 1\}$. Finally, after observing the effort, player 2 decides to hire, $h=1$, or not, $h=0$. The payoffs are

$$u_1(\theta, e, h) = h - c_\theta\, e \quad \text{and} \quad u_2(\theta, e, h) = h\,(2\,\mathbb{I}_{\theta=H} - 1)\ ,$$

where $1 < c_H < c_L < 1/\Delta$. For simplicity, we assume $1/c_L \notin \mathcal{E}$, and we let $\bar{e}$ be the smallest element of $\mathcal{E}$ bigger than $1/c_L$.

We consider the following ℓ-tremble $\tilde{\lambda}$. The low and high types tremble to all actions with likelihoods $(2,1)$ and $(1,1)$, respectively; that is, $\tilde{\lambda}(e|L) = (2,1)$ and $\tilde{\lambda}(e|H) = (1,1)$ for all e. Player 2 trembles to all actions with likelihood $(1,1)$; that is, $\tilde{\lambda}(h|e) = (1,1)$ for all e and h.

Fix some ℓ-equilibrium λ for $\tilde{\lambda}$. Let $e_+ < e$ be the highest effort such that $\lambda(e|L) > \tilde{\lambda}(e|L)$. It must then be that $\lambda(h=1|e) = (0,1)$ for all $e > e_+$. Since, for type L, choosing effort 0 strictly dominates choosing any effort $e \geq \bar{e}$, we have that $e_+ < \bar{e}$. It must also be that

$$\lambda(h|e_+) - c_L\, e_+ \geq 1 - (e_+ + \Delta)\, c_L \quad \Rightarrow \quad c_L \geq \tfrac{1 - \lambda(h|e_+)}{\Delta}\ .$$

By the usual single-crossing property, we have that $\lambda(e|H) = \tilde{\lambda}(e|H)$ for all $e < e_+$. Also, type H is willing to choose e_+ only if $c_H \leq \frac{1-\lambda(h|\bar{e})}{\Delta}$. Since $c_H < c_L$, we have that the high type chooses $e_+ + \Delta$ for sure. Finally, since type L has to assign positive (i.e., no infinitesimal) likelihood to at least one effort below e_+, and since player 2 chooses $h = 0$ if the posterior about the type being H is 0, type L only chooses $e = 0$ with positive likelihood (hence, with probability 1). It then follows that

$$(\lambda(e|L), \lambda(e|H), \lambda(h=1|e)) = \begin{cases} ((0,1),(1,1),(1,1)) & \text{if } e=0, \\ ((1,1),(1,1),(0, e\, c_L)) & \text{if } 0<e<e_+, \\ ((2,1),(0,1),(0,1)) & \text{if } e=e_+, \\ ((2,1),(1,1),(0,1)) & \text{if } e>e_+. \end{cases}$$

Hence, there is a unique ℓ-equilibrium for $\tilde{\lambda}$. This implies that its outcome (which is the Riley outcome), is the unique ℓ-stable outcome, and hence it is stable as well.

7 Conclusions

This paper provides three contributions. The first is developing a new language to analyze limit-based refinements in extensive form games. The ℓ-numbers are simple, 2-dimensional objects, with simple elementary operations. When used to analyze strategies, they permit easily computing relative likelihoods of zero-probability histories.

The second contribution is to use ℓ-numbers to obtain a straightforward characterization of the set sequential equilibria: it coincides with the set of assessments generated by sequentially optimal ℓ-strategy profiles. An implication is that simple sequences are sufficient to generate consistent assessments.

The final contribution is to provide new equilibrium concepts using the language of ℓ-numbers. Most saliently, we define ℓ-stable outcomes, which are the natural analogous of stable outcomes using ℓ-numbers. ℓ-Stable outcomes always exist, are easier to compute, and satisfy desirable properties. We argue that, since an outcome is stable when it is the unique ℓ-stable outcome, our analysis provides a method to obtain stable outcomes in practice, which is illustrated through some examples.

References

Banks, J. S., and J. Sobel, 1987, “Equilibrium selection in signaling games,” *Econometrica: Journal of the Econometric Society*, pp. 647–661.

Battigalli, P., 1996, “Strategic independence and perfect Bayesian equilibria,” *Journal of Economic Theory*, 70(1), 201–234.

Blume, L., A. Brandenburger, and E. Dekel, 1991, “Lexicographic probabilities and equilibrium refinements,” *Econometrica: Journal of the Econometric Society*, pp. 81–98.

Cho, I.-K., and D. M. Kreps, 1987, “Signaling games and stable equilibria,” *The Quarterly Journal of Economics*, 102(2), 179–221.

Fudenberg, D., D. M. Kreps, and D. K. Levine, 1988, “On the robustness of equilibrium refinements,” *Journal of Economic Theory*, 44(2), 354–380.

Fudenberg, D., and J. Tirole, 1991a, *Game theory*. MIT press.

——— , 1991b, “Perfect Bayesian equilibrium and sequential equilibrium,” *Journal of Economic Theory*, 53(2), 236–260.

Govindan, S., and T. Klumpp, 2003, “Perfect equilibrium and lexicographic beliefs,” *International Journal of Game Theory*, 31(2), 229–243.

Jackson, M. O., T. Rodriguez-Barraquer, and X. Tan, 2012, “Epsilon-equilibria of perturbed games,” *Games and Economic Behavior*, 75(1), 198–216.

Kohlberg, E., and J.-F. Mertens, 1986, “On the strategic stability of equilibria,” *Econometrica*, 54(5), 1003–1037.

Kohlberg, E., and P. J. Reny, 1997, “Independence on relative probability spaces and consistent assessments in game trees,” *Journal of Economic Theory*, 75(2), 280–313.

Kreps, D. M., and R. Wilson, 1982, “Sequential equilibria,” *Econometrica: Journal of the Econometric Society*, pp. 863–894.

Mailath, G. J., A. Postlewaite, and L. Samuelson, 2005, “Contemporaneous perfect epsilon-equilibria,” *Games and Economic Behavior*, 53(1), 126–140.

Mailath, G. J., L. Samuelson, and J. M. Swinkels, 1997, “How proper is sequential equilibrium?,” *Games and Economic Behavior*, 18(2), 193–218.

Myerson, R. B., 1986, “Multistage games with communication,” *Econometrica: Journal of the Econometric Society*, pp. 323–358.

Nash, J., 1951, "Non-cooperative games," *Annals of mathematics*, pp. 286–295.

Osborne, M. J., and A. Rubinstein, 1994, *A course in game theory.* MIT press.

Riley, J. G., 1979, "Informational equilibrium," *Econometrica: Journal of the Econometric Society*, pp. 331–359.

Robinson, A., 2016, *Non-standard analysis.* Princeton University Press.

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

Spence, M., 1973, "Job market signaling," *The Quarterly Journal of Economics*, 87(3), 355–374.

Takahashi, S., and O. Tercieux, 2020, "Robust equilibrium outcomes in sequential games under almost common certainty of payoffs," *Journal of Economic Theory*, 188, 105068.

A Omitted Proofs

A.1 Proofs of the Results in Section 3

Proof of Proposition 3.1

Proof. **"Only if" part:** We begin with a useful representation lemma, that shows that any finite collection of sequences on $(0,1]$ can be represented as the product of some basic sequences.

Lemma A.1 (Representation)**.** *Let $(\sigma_n : A \to (0,1])_n$ be a sequence. Then, there are a strictly increasing sequence $(j_n \in \mathbb{N})_n$ and a sequence $((q_n^1,...,q_n^K) \in \mathbb{R}_{++}^K)_n$, for some $K \in \{0,...,|A|\}$, such that*

1. $\lim_{n\to\infty} q_{j_n}^1 = 0$ *and* $\lim_{n\to\infty} q_{j_n}^k/(q_{j_n}^{k-1})^\gamma = 0$ *for all* $\gamma \in \mathbb{R}$ *and* $k = 2,...,K$.
2. *For each* $a \in A$ *there is a unique* $\alpha(a) \equiv (\alpha^1(a),...,\alpha^K(a)) \in \mathbb{R}^K$ *such that*

$$\lim_{n\to\infty} \frac{\sigma_{j_n}(a)}{\prod_{k=1}^K (q_{j_n}^k)^{\alpha^k(a)}} \in \mathbb{R}_{++} \ .$$

Proof. We proceed by induction over the number of actions. If there is one action (i.e., $A = \{a\}$) the result is clear: we let $(j_n)_n$ be strictly increasing and such that $(\sigma_{j_n}(a))_n$ is convergent to some $\sigma(a) \in [0,1]$. If $\sigma(a) > 0$ then $K \equiv 0$, and if $\sigma(a) > 0$ then $K \equiv 1$ and $(q_n^1)_n \equiv (\sigma_n(a))_n$. Assume then that $|A| > 1$ and that the result is true for $A \backslash \{\hat{a}\}$, for some $\hat{a} \in A$, and we prove the result is true for A. We use $(j_n)_n$ and $(q_n^1,...,q_n^K)_n$ to denote the sequences representing $(\sigma_n : A\backslash\{\hat{a}\} \to (0,1])_n$.

We propose the following algorithm, which follows steps $k = K,...,1$ in descending order, from K to when the algorithm stops. In each step, the algorithm takes $(j_n^k)_n$ and $(\tilde{q}_n^k \in \mathbb{R}_{++})_n$ from the previous step, where $\lim_{n\to\infty} \tilde{q}_{j_n^k}^k$ exists (in $\overline{\mathbb{R}}_+$). We initialize (for $k = K$) $(j_n^K)_n$ to be a subsequence of $(j_n)_n$ such that $(\sigma_{j_n^K}(\hat{a}))_n$ is convergent, and $(\tilde{q}_n^K)_n \equiv (\sigma_n(\hat{a}))_n$. The output of the algorithm is a vector $\alpha_* \in \mathbb{R}^K$, two sequences $(j_n^*)_n$ and $(\tilde{q}_n^0,...,\tilde{q}_n^K)_n$, and a value $k^* \in \{0,...,K\}$ indicating the step where the algorithm stopped.

Step 1 (a) If $\liminf_{n\to\infty} \tilde{q}_{j_n^k}^k/(q_{j_n^k}^k)^\gamma = 0$ for all $\gamma \in \mathbb{R}$, then we set $(j_n^*)_n$ to be the subsequence of $(j_n^k)_n$, defined as follows:

i. $j_1^* \equiv j_1^k$.

ii. $j_2^* \equiv \min\{j_n^k > j_1^* | \tilde{q}_{j_n^k}^k/(q_{j_n^k}^k)^2 < 1/2\}$ (exists since $\liminf_{n\to\infty} \tilde{q}_{j_n^k}^k/(q_{j_n^k}^k)^2 = 0$).

iii. ...

iv. $j_m^* \equiv \min\{j_n^k > j_{m-1}^* | \tilde{q}_{j_n^*}^k/(q_{j_n^k}^k)^m < 1/m\}$.

v. ...

Since $\lim_{n\to\infty} q^k_{j^k_n} = 0$ and $\lim_{n\to\infty} \tilde{q}_{j^k_n}/(q^k_{j^k_n})^n = 0$, it is clear that $\lim_{n\to\infty} \tilde{q}^k_{j^*_n}/(q^k_{j^*_n})^\gamma = 0$ for all $\gamma \in \mathbb{R}$. We then set $k^* \equiv k$; set $(\tilde{q}^{k'}_n)_n \equiv (\tilde{q}^k_n)_n$ and $\hat{\alpha}^{k'}_* \equiv 0$ for all $k' \leq k$; and stop.

(b) If the previous case fails, and $\liminf_{n\to\infty} (\tilde{q}^k_{j^k_n})^{-1}/(q^k_{j^k_n})^\gamma = 0$ for all $\gamma \in \mathbb{R}$, then set $(j^*_n)_n$ to be a subsequence of $(j^k_n)_n$ such that $\lim_{n\to\infty} (\tilde{q}^k_{j^*_n})^{-1}/(q^k_{j^*_n})^\gamma = 0$ for all $\gamma \in \mathbb{R}$ (which exists by the previous argument); set $k^* \equiv k$; set $(\tilde{q}^{k'}_n)_n \equiv (\tilde{q}^k_n)_n$ and $\hat{\alpha}^{k'}_* \equiv 0$ for all $k' \leq k$; and stop.

(c) If the previous cases fail and $\liminf_{n\to\infty} q^k_{j^k_n}/(\tilde{q}^k_{j^k_n})^\gamma = 0$ for all $\gamma \in \mathbb{R}$, then $(j^{k-1}_n)_n$ is set to be a subsequence of $(j^k_n)_n$ such that $\lim_{n\to\infty} q^k_{j^{k-1}_n}/(\tilde{q}^k_{j^{k-1}_n})^\gamma = 0$ for all $\gamma \in \mathbb{R}$ (which again exists by the previous argument), $(\tilde{q}^{k-1}_n)_n = (\tilde{q}^k_n)_n$, and $\hat{\alpha}^k_* \equiv 0$; and go to Step 2.

(d) If the previous cases fail, we proceed as follows. Note that it must be that $\lim_{n\to\infty} \tilde{q}^k_{j^k_n} \in \{0, +\infty\}$ (since, otherwise, we would have $\lim_{n\to\infty} q^k_{j^k_n}/(\tilde{q}^k_{j^k_n})^\gamma = 0$ for all $\gamma \in \mathbb{R}$). There are then two sub-cases:

   i. Assume first $\lim_{n\to\infty} \tilde{q}^k_{j_n} = 0$. Since case (a) fails, there is some $\gamma \in \mathbb{R}_{++}$ such that

$$\liminf_{n\to\infty} \frac{\tilde{q}^k_{j^k_n}}{(q^k_{j^k_n})^\gamma} > 0 \ . \tag{A.1}$$

   Let $\gamma_1 \geq 0$ be the infimum of the set of values γ such that the previous inequality holds. Note that, for all $\gamma > \gamma_1$, the left-hand side of (A.1) is $+\infty$; and, if $\gamma < \gamma_1$, the left-hand side of (A.1) is 0. We let $(j^{k-1}_n)_n$ be such that $\tilde{q}^k_{j^{k-1}_n}/(q^k_{j^{k-1}_n})^{\gamma_1}$ tends to some limit in $\overline{\mathbb{R}}_+$. Finally, let $\tilde{q}^{k-1}_n \equiv \tilde{q}^k_n/(q^k_n)^{\gamma_1}$, and $\alpha^k_* = \gamma_1$; and go to Step 2.

   We further prove, for future use, that $\gamma_1 > 0$. To see this, assume for the sake of contradiction that $\gamma_1 = 0$. Then we have that, for all $\gamma > 0$,

$$+\infty = \liminf_{n\to\infty} \frac{\tilde{q}^k_{j^k_n}}{(q^k_{j^k_n})^\gamma} = \left(\limsup_{n\to\infty} \frac{q^k_{j^k_n}}{(\tilde{q}^k_{j^k_n})^{1/\gamma}} \right)^{-\gamma} .$$

   That is, defining $\gamma' \equiv 1/\gamma$, we have that for all $\gamma' > 0$

$$0 = \limsup_{n\to\infty} \frac{q^k_{j^k_n}}{(\tilde{q}^k_{j^k_n})^{\gamma'}} \geq \liminf_{n\to\infty} \frac{q^k_{j^k_n}}{(\tilde{q}^k_{j^k_n})^{\gamma'}} \geq 0 \quad \Rightarrow \quad \liminf_{n\to\infty} \frac{q^k_{j^k_n}}{(\tilde{q}^k_{j^k_n})^{\gamma'}} = 0 \ .$$

   Since $\liminf_{n\to\infty} q^k_{j^k_n}/(\tilde{q}^k_{j^k_n})^{\gamma'} = 0$ for all $\gamma' \leq 0$ (because both $q^k_{j^k_n}$ and $\tilde{q}^k_{j^k_n}$ tend to 0 as $n \to 0$), we reach a contradiction with the assumption that (c) does not hold.

   Note finally, also for future use, that using similar arguments and that case (c) fails, we have that there exists some $\gamma_2 > 0$ such that $\liminf_{n\to\infty} q^k_{j^k_n}/(\tilde{q}^k_{j^k_n})^\gamma =$

$+\infty$ for all $\gamma > \gamma_2$ and $\liminf_{n\to\infty} q^k_{j^k_n}/(\tilde{q}^k_{j^k_n})^\gamma = 0$ for all $\gamma < \gamma_2$. For all $\gamma > \gamma_1$ we have that

$$+\infty = \liminf_{n\to\infty} \frac{\tilde{q}^k_{j^k_n}}{(q^k_{j^k_n})^\gamma} = \left(\limsup_{n\to\infty} \frac{q^k_{j^k_n}}{(\tilde{q}^k_{j^k_n})^{1/\gamma}} \right)^{-\gamma} \le \left(\liminf_{n\to\infty} \frac{q^k_{j^k_n}}{(\tilde{q}^k_{j^k_n})^{1/\gamma}} \right)^{-\gamma}.$$

Hence, we have $\gamma > \gamma_1$ implies $1/\gamma < \gamma_2$; that is, $\gamma_1 \ge 1/\gamma_2$. Similarly, if $\gamma < \gamma_1$, we have

$$0 = \liminf_{n\to\infty} \frac{\tilde{q}^k_{j^k_n}}{(q^k_{j^k_n})^\gamma} = \left(\liminf_{n\to\infty} \frac{q^k_{j^k_n}}{(\tilde{q}^k_{j^k_n})^{1/\gamma}} \right)^{-\gamma}.$$

Now we have that $\gamma < \gamma_1$ implies $1/\gamma > \gamma_2$, that is, $\gamma_1 \le 1/\gamma_2$. Overall, we have $\gamma_1 = 1/\gamma_2$.

ii. Assume now $\lim_{n\to\infty} \tilde{q}^k_{j^k_n} = +\infty$. We proceed analogously, where now γ_1 is the *supremum* of the set of values γ such that inequality (A.1) holds. The rest holds equivalently.

Step 2 If $k > 1$ then replace k by $k-1$ and go to Step 1. If $k = 1$ then set $(j^*_n)_n \equiv (j^0_n)_n$ and $k^* \equiv 0$, and stop.

(End of the algorithm, the proof of Proposition 3.1 continues.)

Define $(q^*_n)_n \equiv (\tilde{q}^0_n)_n$. We note first that

$$q^*_n = \tilde{q}^0_n = \frac{\tilde{q}^1_n}{(q^1_n)^{\alpha^1_*}} = \frac{1}{(q^1_n)^{\alpha^1_*}} \frac{\tilde{q}^2_n}{(q^2_n)^{\alpha^2_*}} = \dots = \frac{\sigma_n(\hat{a})}{\prod_{k=1}^K (q^k_{j_n})^{\alpha^k_*}}. \tag{A.2}$$

We note also that, for all $k > k^*$, it must be that

$$\lim_{n\to\infty} \frac{q^k_{j^*_n}}{(\tilde{q}^{k-1}_{j^*_n})^\gamma} = 0 \quad \text{for all } \gamma \in \mathbb{R}. \tag{A.3}$$

Indeed, $k > k^*$ means that both cases 1(a) and 1(b) in the algorithm do not hold for k. If case 1(c) holds for k, then it is clear that $\lim_{n\to\infty} q^k_{j^*_n}/(\tilde{q}^{k-1}_{j^*_n})^\gamma = 0$ for all $\gamma \in \mathbb{R}$. Assume then case 1(d) holds for k, and that $\lim_{n\to\infty} \tilde{q}^k_{j^*_n} = 0$ (the case where $\lim_{n\to\infty} \tilde{q}^k_{j^*_n} = +\infty$ is analogous). Assume then, for the sake of contradiction, that there is some $\gamma \in \mathbb{R}$ such that

$$0 < \limsup_{n\to\infty} \frac{q^k_{j^*_n}}{(\tilde{q}^{k-1}_{j^*_n})^\gamma} = \limsup_{n\to\infty} \frac{(q^k_{j^*_n})^{1+\gamma\gamma_1}}{(\tilde{q}^k_{j^*_n})^\gamma}, \tag{A.4}$$

where we used that $\tilde{q}^{k-1}_n \equiv \tilde{q}^k_n/(q^k_n)^{\gamma_1}$ in case 1(d). If $1 + \gamma\gamma_1 = 0$, then it must be that $\gamma < 0$ (since $\gamma_1 > 0$), but then equation (A.4) does not hold because both $q^k_{j^*_n}$ and

$(\tilde{q}^{k-1}_{j^*_n})^{-\gamma}$ tend to 0. Alternatively, if $1+\gamma\,\gamma_1>0$, then it must be that $\gamma>0$ for equation (A.4) to hold, and hence

$$0<\limsup_{n\to\infty}\frac{(q^k_{j^*_n})^{1+\gamma\,\gamma_1}}{(\tilde{q}^k_{j^*_n})^{\gamma}}=\left(\liminf_{n\to\infty}\frac{\tilde{q}^k_{j^*_n}}{(q^k_{j^*_n})^{(1+\gamma\,\gamma_1)/\gamma}}\right)^{-\gamma}.$$

By the definition of γ_1, this implies that $(1+\gamma\,\gamma_1)/\gamma\le\gamma_1$, which is a contradiction. Finally, if $1+\gamma\,\gamma_1<0$, then it must be that $\gamma<0$ (since $\gamma_1>0$), and hence we have

$$0<\limsup_{n\to\infty}\frac{(q^k_{j^*_n})^{1+\gamma\,\gamma_1}}{(\tilde{q}^k_{j^*_n})^{\gamma}}=\left(\limsup_{n\to\infty}\frac{\tilde{q}^k_{j^*_n}}{(q^k_{j^*_n})^{(1+\gamma\,\gamma_1)/\gamma}}\right)^{-\gamma}=\left(\liminf_{n\to\infty}\frac{q^k_{j^*_n}}{(\tilde{q}^k_{j^*_n})^{\gamma/(1+\gamma\,\gamma_1)}}\right)^{1+\gamma\,\gamma_1}.$$

By the definition of γ_2 (in Step 1(d) of the algorithm), we have that $\gamma/(1+\gamma\,\gamma_1)<\gamma_2=1/\gamma_1$, which is again a contradiction.

To conclude the proof, note that there are three possibilities:

1. If $\lim_{n\to\infty} q^*_{j^*_n}\in\mathbb{R}_{++}$ (and so $k^*=0$) then we have that $(j^*_n)_n$ and $(q^1_n,\ldots,q^K_n)_n$ are the desired sequence (by (A.2) we have that the second property of the statement holds for $\hat{a}$ for $\alpha(\hat{a})=\alpha_*$).

2. If $\lim_{n\to\infty} q^*_{j^*_n}=0$ then $(j^*_n)_n$ and

$$(q^1_n,\ldots.,q^{k^*}_n,q^*_n,q^{k^*+1}_n,\ldots,q^K_n)_n$$

are the desired sequences. Indeed, in this case, the algorithm ends in one of the following two cases. In the first cases, the algorithm ends because $k^*=0$. In this case, since $(q^*_n)_n=(\tilde{q}^0_n)_n$, we have $\lim_{n\to\infty} q^1_{j^*_n}/(q^*_{j^*_n})^{\gamma}=0$ for all $\gamma\in\mathbb{R}$ (by equation (A.3) at $k=1$). In the second case, the algorithm ends at Step 1(a) for some $k^*>0$. Nevertheless, in this case, for all $\gamma\in\mathbb{R}$,

$$\lim_{n\to\infty}\frac{q^*_{j^*_n}}{(q^{k^*}_{j^*_n})^{\gamma}}=0\ \text{ and, if } k^*<K,\ \lim_{n\to\infty}\frac{q^{k^*+1}_{j^*_n}}{(q^*_{j^*_n})^{\gamma}}=0\ .$$

The first equality holds from the algorithm ends at Step 1(a) for k^*, while the second equality holds because $(q^*_n)_n=(\tilde{q}^{k^*}_n)_n$ (and using equation (A.3) again, now at $k=k^*+1$).

3. If $\lim_{n\to\infty}(q^*_{j^*_n})^{-1}=0$ then, proceeding as in the previous case, it is easy to see that $(j^*_n)_n$ and

$$(q^1_n,\ldots.,q^{k^*}_n,(q^*_n)^{-1},q^{k^*+1}_n,\ldots,q^K_n)$$

are the desired sequences.

□

(End of proof of Lemma A.1, proof of Proposition 3.1 continues.)

Fix a consistent assessment (σ, μ) and a sequence of fully mixed strategies $(\sigma_n)_n$ with corresponding beliefs $(\mu_n)_n$ (obtained using Bayes' rule) converging to it. Let $(j_n)_n$ and $(q_n^1,\dots,q_n^K)_n$ be two sequences with the properties in the statement of Lemma A.1 for σ_n. We then have that, for each action $a \in A$, there exist some unique $\alpha^0(a) \in \mathbb{R}_{++}$ and $\alpha(a) \in \mathbb{R}$ such that

$$\lim_{n\to\infty} \frac{\sigma_{j_n}(a)}{\alpha^0(a)\, \prod_{k=1}^K (q_{j_n}^k)^{\alpha^k(a)}} = 1 \ .$$

Similarly, it follows that, for each history $h = (a_1, \dots, a_J) \in H$, we have

$$\lim_{n\to\infty} \frac{\Pr(h|\sigma_{j_n})}{\left(\prod_{j=1}^J \alpha^0(a_j)\right) \prod_{k=1}^K (q_{j_n}^k)^{\alpha^k(h)}} = 1 \ ,$$

where $\alpha(h) \equiv \sum_{j=1}^J \alpha(a_j)$.

For each $M \in \mathbb{R}_{++}$, we define

$$\hat{q}_{M,n} \equiv \left(q_n^1/M, q_n^2/M^2, \dots., q_n^K/M^K\right)$$

and

$$\hat{\alpha}_M(a) \equiv \left(M\, \alpha^1(a), M^2\, \alpha^2(a), \dots., M^K\, \alpha^K(a)\right) \ .$$

For each M, we define the ℓ-strategy profile λ_M assigning, to each action $a \in A$, the value

$$\lambda_M(a) \equiv \left(\textstyle\sum_{k=1}^K \hat{\alpha}_{M,k}(a), \alpha^0(a)\right) \ . \qquad \text{(A.5)}$$

We now argue that, if M is large enough, then for all histories $h, h' \in H$,

$$\lim_{n\to\infty} \frac{\sigma_{j_n}(h)}{\sigma_{j_n}(h')} = \begin{cases} \lambda_M^{\mathrm{p}}(h)/\lambda_M^{\mathrm{p}}(h') & \text{if } \lambda_M^{\mathrm{r}}(h) = \lambda_M^{\mathrm{r}}(h'), \\ 0 & \text{if } \lambda_M^{\mathrm{r}}(h) < \lambda_M^{\mathrm{r}}(h'), \\ \infty & \text{if } \lambda_M^{\mathrm{r}}(h) > \lambda_M^{\mathrm{r}}(h'). \end{cases} \qquad \text{(A.6)}$$

To see this, we first define $\overline{k}(\hat{\alpha}) \equiv \max(\{0\} \cup \{k | \hat{\alpha}^k \neq 0\})$ for each $\hat{\alpha} \in \mathbb{R}^K$. We then note that, if M is large enough, the left-hand side of expression (A.6) is $\alpha^0(h)/\alpha^0(h')$ if and only if $\hat{\alpha}_M(h) = \hat{\alpha}_M(h')$, in which case $\lambda_M^{\mathrm{r}}(h) = \lambda_M^{\mathrm{r}}(h')$. If, instead, $\lambda_M^{\mathrm{r}}(h) < \lambda_M^{\mathrm{r}}(h')$ for large M then it must be that

$$\overline{k}(\hat{\alpha}_M(h)) < \overline{k}(\hat{\alpha}_M(h')), \text{ or } \overline{k}(\hat{\alpha}_M(h)) = \overline{k}(\hat{\alpha}_M(h')) \text{ and } \alpha_M^{\overline{k}(\hat{\alpha}_M(h))}(h) < \alpha_M^{\overline{k}(\hat{\alpha}_M(h'))}(h'),$$

which by the second property in Lemma A.1 implies the left-hand side of expression (A.6) is 0. Analogous arguments imply that, if $\lambda_M^{\mathrm{r}}(h) > \lambda_M^{\mathrm{r}}(h')$ and M is large enough, then the left-hand side of expression (A.6) is $+\infty$.

The previous arguments imply that, if M is large enough then, for two histories $h, h' \in I$, equation (A.6) holds, which implies that (σ, μ) is generated by λ_M, and hence the "Only if" part of the statement is proven.

"If" part: The arguments in the main text show that, if λ is an ℓ-strategy profile, then the sequence defined in equation (3.1) supports $(\sigma^\lambda, \mu^\lambda)$, hence the proof is completed. □

Proof of the Proposition 3.2

Proof. Fix first some ℓ-equilibrium λ. The standard "one-period deviation principle" implies that $(\sigma^\lambda, \mu^\lambda)$ is a sequential equilibrium if and only if for all $I \in \mathcal{I}$ and $a, a' \in A^I$ such that $\sigma^\lambda(a) > 0$ we have $u(a|\sigma^\lambda, \mu^\lambda) \geq u(a'|\sigma^\lambda, \mu^\lambda)$. Note also that, for all $a \in A$, $\sigma^\lambda(a) > 0$ if only if $\lambda^{\mathrm{r}}(a) = 0$, and also $u(a|\sigma^\lambda, \mu^\lambda) = u(a|\lambda)$. This implies that $(\sigma^\lambda, \mu^\lambda)$ sequential equilibrium.

Now fix some sequential equilibrium (σ, μ). Let λ be such that $(\sigma^\lambda, \mu^\lambda) = (\sigma, \mu)$ (note that, by Proposition 3.1, λ exists). That λ is an ℓ-equilibrium follows from the fact that, for any $I \in \mathcal{I}$ and $a \in I$, we have $\lambda^{\mathrm{r}}(a) = 0$ if and only if $\sigma^\lambda(a) > 0$, that is, if and only if $u(a|\sigma^\lambda) \geq \max_{a' \in A^I} u(a'|\sigma^\lambda)$, which implies $u(a|\lambda) \geq \max_{a' \in A^I} u(a'|\lambda)$. □

A.2 Proofs of the Results in Section 4

Proof of Proposition 4.1

Proof. The proof follows from the arguments in the main text. □

Proof of Proposition 4.2

Proof. **"Only if" part:** Assume λ is an ℓ-equilibrium for $\tilde{\lambda}$. Define $(\sigma_n)_n$ as in (3.1). Note that $\sigma_n(a) = \eta_n^{\tilde{\lambda}}(a)$ for all n whenever $\lambda^{\mathrm{r}}(a) > 0$ (recall Definition 4.5). For each n, we let $\varepsilon_n \geq 0$ be the minimum epsilon such that $\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta_n^{\tilde{\lambda}})$. Given that $(\sigma^\lambda, \mu^\lambda)$ is a sequential equilibrium supported by $(\sigma_n)_n$, it follows that $\varepsilon_n \to 0$, and hence $(\sigma^\lambda, \mu^\lambda)$ is a perfect equilibrium for $(\eta_n^{\tilde{\lambda}})_n$.

"If" part: Assume there is a sequence $(\sigma_n)_n$ supporting $(\sigma^\lambda, \mu^\lambda)$, and two sequences $(\varepsilon_n)_n \searrow 0$ and $(\eta_n)_n$, such that $\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta_n)$ for all n. It is clear that $(\sigma^\lambda, \mu^\lambda)$ is a sequential equilibrium. Therefore, by Propositions 3.2 and 4.1, λ is an ℓ-equilibrium for some ℓ-tremble. □

Proof of Proposition 4.3

Proof. We divide the proof into four steps:

Step 1. An auxiliary result. We begin by proving a useful result:

Lemma A.2. *Let λ be an ℓ-equilibrium for $\tilde{\lambda}$. Define, for each $\kappa, \gamma > 0$,*

$$(\lambda^r_{\kappa,\gamma}(a), \lambda^p_{\kappa,\gamma}(a)) \equiv (\kappa\, \lambda^r(a), \gamma^{\lambda^r(a)}\, \lambda^p(a)) \quad \forall a \in A$$

and

$$(\tilde{\lambda^r}_{\kappa,\gamma}(a), \tilde{\lambda^p}_{\kappa,\gamma}(a)) \equiv (\kappa\, \tilde{\lambda^r}(a), \gamma^{\tilde{\lambda^r}(a)}\, \tilde{\lambda^p}(a)) \quad \forall a \in A\ .$$

Then, $\lambda_{\kappa,\gamma}$ is an ℓ-equilibrium for $\tilde{\lambda}_{\kappa,\gamma}$.

Proof. We first note that $\lambda_{\kappa,\gamma}$ is an ℓ-strategy profile. The reason is that conditions 1 and 2 in Definition 3.1 are trivially satisfied since, whenever $\lambda(a) = 0$, we have $\lambda_{\kappa,\gamma}(a) = \lambda(a)$. It is also clear that $\lambda_{\kappa,\gamma}$ supersedes $\tilde{\lambda}_{\kappa,\gamma}$ (in the sense of Definition 4.2).

It is only left to prove that $\lambda_{\kappa,\gamma}(a) > \tilde{\lambda}_{\kappa,\gamma}(a)$ for $a \in A$ only if $u(a|\lambda_{\kappa,\gamma}) \geq u(a'|\lambda_{\kappa,\gamma})$ for all $a' \in A^{I^a}$. Since, for all $a \in A$, $\lambda_{\kappa,\gamma}(a) > \tilde{\lambda}_{\kappa,\gamma}(a)$ if and only $\lambda(a) > \tilde{\lambda}(a)$, it is enough to show that $(\sigma^{\lambda_{\kappa,\gamma}}, \mu^{\lambda_{\kappa,\gamma}}) = (\sigma^\lambda, \mu^\lambda)$. It is clear that $\sigma^{\lambda_{\kappa,\gamma}} = \sigma^\lambda$. Note now that, for a history $h \equiv (a_j)_{j=1}^J$,

$$\begin{aligned} \lambda_{\kappa,\gamma}(h) &= \left(\textstyle\sum_{j=1}^J \lambda^{\mathrm{r}}_{\kappa,\gamma}(a_j), \prod_{j=1}^J \lambda^{\mathrm{p}}_{\kappa,\gamma}(a_j) \right) \\ &= \left(\kappa \textstyle\sum_{j=1}^J \lambda^{\mathrm{r}}(a_j), \gamma^{\sum_{j=1}^J \lambda^{\mathrm{r}}(a_j)} \prod_{j=1}^J \lambda^{\mathrm{p}}(a_j) \right) \\ &= \left(\kappa\, \lambda^{\mathrm{r}}(h), \gamma^{\lambda^{\mathrm{r}}(h)}\, \lambda^{\mathrm{p}}(h) \right)\ . \end{aligned}$$

It is then clear that $\mu^{\lambda_{\kappa,\gamma}} = \mu^\lambda$. It then follows that $\lambda_{\kappa,\gamma}$ is an ℓ-equilibrium for $\tilde{\lambda}_{\kappa,\gamma}$. □

(End of the proof of Lemma A.2. Proof of Proposition 4.3 continues.)

Step 2. "Only if" part of the second statement. We now prove that, if λ is an ℓ-equilibrium for $\tilde{\lambda}$, then there is a sequence $(\sigma_n)_n$ supporting $(\sigma^\lambda, \mu^\lambda)$, and a sequence $(\varepsilon_n)_n \searrow 0$ such that $\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta_n^{\tilde{\lambda}})$ for all n. For each $a \in A$, we define $\sigma_n(a)$ as in equation (3.1). Fix some $\varepsilon > 0$. It is clear that σ_n is a strategy profile, and that it satisfies the first condition of the definition of ε-equilibrium for η_n (Definition 4.3). Now, let $a \in A$ be such that $\sigma_n(a) > \eta_n(a)$. This only happens if $\lambda(a) > \tilde{\lambda}(a)$, which requires that

$$u(a|\lambda) \geq u(a'|\lambda) \quad \text{for all } a' \in A^{I^a},$$

since λ is an ℓ-equilibrium for $\tilde{\lambda}$. Since $u(a|\lambda) = \lim_{n\to\infty} u(a|\sigma_n)$, we have that

$$\lim_{n\to\infty} \left(u(a|\sigma_n) - u(a'|\sigma_n) \right) \geq 0\ .$$

Therefore, for each $a' \in A^I$, there exists some $\overline{n}_{a,a'}$ such that

$$u(a|\sigma_n) \geq u(a'|\sigma_n) - \varepsilon \quad \text{for all } n > \overline{n}_{a,a'} \ .$$

We let $\overline{n}$ be the maximum $\overline{n}_{a,a'}$ across all I, $a \in I$ with $\lambda(a) > \tilde{\lambda}(a)$, and $a' \in I$. It is then clear that, for all $n > \overline{n}$, σ_n is an ε-equilibrium for $\eta_n^{\tilde{\lambda}}$. Standard arguments imply that a sequence $(\varepsilon_n)_n \searrow 0$ such that $\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta_n^{\tilde{\lambda}})$ for all n.

Step 3: "If" part of the second statement. We now fix some assessment (σ, μ) and assume it is a perfect equilibrium for $(\eta_n^{\tilde{\lambda}})_n$ for some $\tilde{\lambda}$. We will show that (σ, μ) is generated by some ℓ-equilibrium for $\tilde{\lambda}$.

Let $(\sigma_n)_n$ be a sequence supporting (σ, μ), and $(\hat{\varepsilon}_n)_n \searrow 0$ be such that $\sigma_n \in \Sigma^*_{\hat{\varepsilon}_n}(\eta_n^{\tilde{\lambda}})$ for all n. Let $(j_n)_n$ and $(q_n^1, ..., q_n^K)_n$ be some sequences satisfying the properties in the statement of Lemma A.1 for $(\sigma_n)_n$. Hence, for any $a \in A$, there are some unique $\alpha_0(a) \in \mathbb{R}_{++}$ and $\alpha(a) \in \mathbb{R}$, such that

$$\lim_{n\to\infty} \frac{\sigma_{j_n}(a)}{\alpha^0(a) \prod_{k=1}^K (q_{j_n}^k)^{\alpha^k(a)}} = 1 \ .$$

The arguments after Lemma A.1 show that, for any M large enough, λ_M defined in equation (A.5) is such that $(\sigma^{\lambda_M}, \mu^{\lambda_M}) = (\sigma, \mu)$. Nevertheless, there is no guarantee that λ_M is an ℓ-equilibrium for $\tilde{\lambda}$. We now show that λ_M can be conveniently transformed using Lemma A.2 so that it becomes an ℓ-equilibrium for $\tilde{\lambda}$.

Let $A^\dagger$ be the set of actions in $a \in A$ satisfying that $\sigma_n(a) > \eta_n(a)$ for infinitely many n. It is clear that, if $a \in A^\dagger$, then

$$u(a|\sigma) \geq u(a'|\sigma) \ \text{ for all } a' \in A^{I^a}.$$

There are two cases:

1. If $A^\dagger \neq A$ then we proceed as follows. Let $\hat{a}$ be some action in $A \backslash A^\dagger$. Then, for each M, we choose κ_M and γ_M so that $\lambda_{M,\kappa_M,\gamma_M}(\hat{a}) = \tilde{\lambda}(\hat{a})$, where we use Lemma A.2 to define

$$\lambda_{M,\kappa_M,\gamma_M}(a) \equiv \left(\kappa_M \textstyle\sum_{k=1}^K \alpha_k(a)/M^k, (\gamma_M)^{\sum_{k=1}^K \alpha_k(a)/M^k} \alpha^0(a)\right)$$

for all $a \in A$. It then follows that $\lambda_{M,\kappa_M,\gamma_M}(\hat{a}') = \tilde{\lambda}(\hat{a}')$ for all $\hat{a}' \in A \backslash A^\dagger$. Indeed, from Definition 4.5, we have

$$\eta_n(\hat{a}') = \frac{\tilde{\lambda}^{\mathrm{p}}(\hat{a})}{\tilde{\lambda}^{\mathrm{p}}(\hat{a})^{\tilde{\lambda}^{\mathrm{r}}(\hat{a}')/\tilde{\lambda}^{\mathrm{r}}(\hat{a})}} \, \eta_n(\hat{a})^{\tilde{\lambda}^{\mathrm{r}}(\hat{a}')/\tilde{\lambda}^{\mathrm{r}}(\hat{a})} \ .$$

Hence, we have $\alpha(\hat{a}') = (\tilde{\lambda}^{\mathrm{r}}(\hat{a}')/\tilde{\lambda}^{\mathrm{r}}(\hat{a}))\, \alpha(\hat{a})$, and so

$$\lambda^{\mathrm{r}}_{M,\kappa_M,\gamma_M}(\hat{a}') = (\tilde{\lambda}^{\mathrm{r}}(\hat{a}')/\tilde{\lambda}^{\mathrm{r}}(\hat{a}))\, \lambda^{\mathrm{r}}_{M,\kappa_M,\gamma_M}(\hat{a}') \ .$$

Similarly,

$$\lambda^{\mathrm{P}}_{M,\kappa_M,\gamma_M}(\hat{a}') = (\tilde{\lambda}^{\mathrm{P}}(\hat{a}')/\tilde{\lambda}^{\mathrm{P}}(\hat{a}))\, \lambda^{\mathrm{P}}_{M,\kappa_M,\gamma_M}(\hat{a}') \ .$$

Also, by the same arguments as the ones used after Lemma A.1, it must be that $\lambda_{M,\kappa_M,\gamma_M}(a) \geq \tilde{\lambda}(a)$ for all $a \in A$ if M is large enough. It then follows that $\lambda_{M,\kappa_M,\gamma_M}$ is an ℓ-equilibrium for $\tilde{\lambda}$.

2. If $A^{\dagger} = A$ then note that λ_M is a ℓ-equilibrium. Using κ large enough, we have that $\lambda_{M,\kappa,1}(a) \geq \tilde{\lambda}(a)$ for all a. Then, since $u(a|\lambda_{M,\kappa,1}) = u(a'|\lambda_{M,\kappa,1})$ for all $a \in A$ and $a' \in A^{I^a}$, $\lambda_{M,\kappa,1}$ an ℓ-equilibrium for $\tilde{\lambda}$.

Step 4: Proof of existence. We now fix some ℓ-tremble $\tilde{\lambda}$, and we prove an ℓ-equilibrium for $\tilde{\lambda}$ exists. For each $n \in \mathbb{N}$, we take some $\sigma_n \in \Sigma^*_0(\eta^{\tilde{\lambda}}_n)$. It is clear that each σ_n is an ε-equilibrium for all $\varepsilon > 0$. Taking a subsequence if necessary, we assume that the sequence of assessments implied by $(\sigma_n)_n$ converges to some assessment (σ, μ) (that is, $(\sigma_n)_n$ supports (σ, μ)). The proof is then concluded using Step 2 of the proof. □

A.3 Proofs of the Results in Section 5

Proof of Proposition 5.1

See the proof in the main text.

Proof of Corollary 5.2

Proof. The proof is immediate from the definition of ℓ-stable outcome and Proposition 5.1. □

Proof of Proposition 5.2

Proof. **"Only if" part.** Assume first o is an ℓ-stable outcome. Fix some $\tilde{\lambda}$. We want to prove that there are two sequences $(\varepsilon_n)_n \searrow 0$ and $(\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta^{\tilde{\lambda}}_n))_n$ such that $o^{\sigma_n} \to o$.

Let λ be a ℓ-equilibrium for $\tilde{\lambda}$ such that $o^{\lambda} = o$ (which exists since o is ℓ-stable). Let then $(\sigma_n)_n$ defined as in equation (3.1). It is clear that $\sigma_n(a) \geq \eta^{\tilde{\lambda}}_n(a)$ for all n, and that $\sigma_n(a) > \eta^{\tilde{\lambda}}_n(a)$ only if $u(a|\lambda) \geq u(a'|\lambda)$ for all $a' \in A^{I^a}$. It then follows that, since $u(a|\sigma_n) \to u(a|\lambda)$ as $n \to \infty$, there exists a sequence of $(\varepsilon_n)_n \searrow 0$ such that $(\sigma_n \in \Sigma^*_{\varepsilon_n}(\eta^{\tilde{\lambda}}_n))_n$ such that $o^{\sigma_n} \to o$.

"If" part. Assume now that o is such that there is some perfect equilibrium (σ, μ) for $(\eta^{\tilde{\lambda}}_n)_n$ with outcome o, for some $\tilde{\lambda} \in \tilde{\Lambda}$. By Proposition 4.3, there is an ℓ-equilibrium for $\tilde{\lambda}$ such that which generates (σ, μ), and hence has outcome o. The proof is then concluded. □

Proof of Corollary 5.1

Proof. Let o be a stable outcome. Assume, for the sake of contradiction, that o is not ℓ-stable. Let $\tilde{\lambda}$ be such that there is no ℓ-equilibrium for $\tilde{\lambda}$ with outcome o. Let $(\sigma_n \in \Sigma_0^*(\eta_n^{\tilde{\lambda}}))_n$ be a sequence such that the corresponding sequence of outcomes converging to o (note that such sequence exists because o is stable). Taking a subsequence if necessary, assume that $(\sigma_n)_n$ supports some assessment (σ, μ), and note that the outcome of (σ, μ) is o. Since $\sigma_n \in \Sigma_\varepsilon^*(\eta_n^{\tilde{\lambda}})$ for all $\varepsilon > 0$, Proposition 4.3 implies that there exists some ℓ-equilibrium for $\tilde{\lambda}$ generating (σ, μ), and hence with outcome o, which is a contradiction. □

Proof of Proposition 5.3

Proof. Assume that o is the unique ℓ-stable outcome and, for the sake of contradiction, assume that it is not stable. Note that there is no stable outcome different from o, since otherwise such an outcome would also be ℓ-stable by Corollary 5.1. Hence, it must be that o is not a stable outcome.

Let $(\hat{\eta}_n)_n$ be a tremble such that there is no sequence $(\sigma_n \in \Sigma_0^*(\hat{\eta}_n))_n$ with outcomes converging to o (which exists since o is not stable). Let $(u_m)_m$ be a sequence converging to u such that there is a stable outcome $o_{n,m}$ of $G(\hat{\eta}_n, u_m)$ for each n and m. Assume that, for each n, $(o_{n,m})_m$ converges to some outcome o_n, which is an equilibrium outcome of $G(\hat{\eta}_n, u)$. Taking a subsequence if necessary, assume that $(o_n)_n$ converges to some outcome o', which is necessarily different from o. Let $(m_n)_n$ be a sequence such that $(o_{n,m_n})_n$ converges to o'.

Fix now some $\tilde{\lambda} \in \tilde{\Lambda}$ and, for each $j \in \mathbb{N}$, a sequence $(\sigma_{nj} \in \Sigma_0^*(\eta_n^{\tilde{\lambda}}, G(\hat{\eta}_j, u_{m_j})))_n$ with outcomes converging to o_{j,m_j}, where $\Sigma_0^*(\eta_n^{\tilde{\lambda}}, G(\hat{\eta}_j, u_{m_j}))$ is the set of equilibria of the game $G(\hat{\eta}_j, u_{m_j})$ where players tremble according to $\eta_n^{\tilde{\lambda}}$ (note that $(\sigma_{nj})_n$ exists since o_{j,m_j} is a stable outcome of $G(\hat{\eta}_j, u_{m_j})$). We define $(j_n)_n$ recursively as follows:

1. For $n = 1$, let j_1 be the smallest such that $\sigma_{1j_1} \in \Sigma_1^*(\eta_1^{\tilde{\lambda}})$.
2. For each $n > 1$, let $j_n > j_{n-1}$ be such that $\sigma_{nj_n} \in \Sigma_{1/n}^*(\eta_n^{\tilde{\lambda}})$.

Let $\sigma_n \equiv \sigma_{nj_n}$ for all $n \in \mathbb{N}$. It is clear that the sequence of outcomes of $(\sigma_n)_n$ converges to o'. Let finally $(n_k \in \mathbb{N})_k$ be strictly increasing and such that $(\sigma_{n_k})_k$ generates some assessment (σ, μ). An analogous argument as the one used in the proof of Proposition 4.3 implies that there exists some $\lambda \in \Lambda^*(\tilde{\lambda})$ such that $(\sigma^\lambda, \mu^\lambda) = (\sigma, \mu)$, and hence $o^\lambda = o'$. Since the argument can be applied to any $\tilde{\lambda} \in \tilde{\Lambda}$, we conclude that o' is ℓ-stable, which is a contradiction. Hence, o is stable. □

Proof of Proposition 5.4

Proof. The proof is analogous to the proof of Proposition 5.2. □

Proof of Proposition 5.5

Proof. 1. *Forward induction:* The proof is in the main text.

2. *Iterated strict dominance:* The argument is analogous the proof of part 1.

3. *Invariance to reordering simultaneous moves:* Let $I, I' \in \mathcal{I}$ be such that $I' = I \times A^I$. For each terminal history $(a_1, ..., a_J) \in T$, define

$$\mathcal{T}(a_1, ..., a_J) \equiv \begin{cases} (a_1, ..., a_{j+1}, a_j, ..., a_J) & \text{if } a_j \in A^I \text{ for some } j, \\ (a_1, ..., a_J) & \text{otherwise.} \end{cases}$$

Let G' be a game coinciding with G but replacing T by $\mathcal{T}(T)$, and also replacing H, $\mathcal{I}$, and ι accordingly.

We now fix some ℓ-stable outcome o of G, and we will show that the outcome analogous to o in G', denoted $o' \equiv o(\mathcal{T}^{-1})$, is also ℓ-stable. To see this, fix some ℓ-tremble $\tilde{\lambda} \in \tilde{\Lambda}$, and let λ be an ℓ-equilibrium for $\tilde{\lambda}$ with outcome o. We argue that λ is also an ℓ-equilibrium for $\tilde{\lambda}$ in G', and has outcome o'. This result follows from the observation that the optimality requirement in the definition of ℓ-equilibrium for $\tilde{\lambda}$ (second condition in Definition 4.2) is equivalent to require that $\lambda(a) > \tilde{\lambda}(a)$ only if $\hat{u}(a|\lambda) \geq \hat{u}(a'|\lambda)$ for all $a' \in A^{I^a}$, where

$$\hat{u}(a|\lambda) \equiv \sum_{t \in T^a} \frac{\lambda(t)}{\lambda(a)} \, u_{\iota(I^a)}(t) \ ,$$

instead of using $u(a|\lambda)$ defined in equation (3.4). Since $\hat{u}(a|\lambda)$ is independent of whether I and I' are reversed or not, it follows that λ is also an ℓ-equilibrium for $\tilde{\lambda}$ in G'.

□

B Relationship to LPSs and CPSs

In this section, we shed light on the link between, on one hand, ℓ-strategy profilesand (i) lexicographic probability systems (LPSs) and (ii) conditional probability systems (CPSs), as they have also been used to characterize consistent assessments.

B.1 Lexicographic Probability Systems

We first relate ℓ-strategy profiles and LPSs. We follow an approach similar to Govindan and Klumpp (2003); that is, we consider LPSs over each set of individual strategies. Such construction is easier to use than that in Blume, Brandenburger, and Dekel (1991), where LPSs are over the space of strategy profiles, which has not finite representation and requires additional conditions to ensure that players randomize independently.

In a normal-form game, Govindan and Klumpp define an LPS as a finite sequence of mixed strategies. To simplify the exposition, we introduce the concept of *behavior LPS* as a sequence of behavior strategies.

Definition B.1. A *(full-support) behavior LPS (profile)* is a finite sequence of strategy profiles $\hat{\sigma} \equiv (\sigma^j)_{k=0}^{\hat{K}}$, for some $\hat{K} \in \mathbb{Z}_+$, such that, for all $a \in A$, there is some k such that $\sigma^k(a) > 0$. We use $\hat{K}(a)$ to denote $\min\{k | \sigma^k(a) > 0\}$.

As Govindan and Klumpp point out, $(\sigma^k|_{\mathcal{I}})_{k=1}^{\hat{K}}$ can be interpreted as a collection of theories of players $N \backslash \{\iota(I)\}$ about player $i(I)$'s strategy, ordered in decreasing likelihood. Using this interpretation, we can define the assessment generated by some behavior LPS $\hat{\sigma}$ by assigning, to each action a and each history $(a_j)_{j=1}^J$ in some information set I, the following probability and belief:

$$\sigma(a) = \sigma^0(a) \quad \text{and} \quad \mu((a)_{j=1}^J) = \begin{cases} 0 & \text{if } \hat{K}((a_j)_{j=1}^J) > \hat{K}(I), \\ C_I \prod_{j=1}^J \sigma^{\hat{K}(a_j)}(a_j) & \text{if } \hat{K}((a_j)_{j=1}^J) = \hat{K}(I), \end{cases} \tag{B.1}$$

where $\hat{K}((a_j)_{j=1}^J) = \sum_{j=1}^J \hat{K}(a_j)$ and $\hat{K}(I) = \min_{h \in I} \hat{K}(h)$, and where C_I is the constant that keeps $\mu|_I$ a probability distribution. In words, the strategy profile in the assessment generated by the LPS coincides with the primary theory about each player's strategy, and the belief at a given information set is derived by the most likely theory from the LPS conditional on the information set being reached.

The following result illustrates the relationship between ℓ-strategy profiles and (behavior) LPSs:

Proposition B.1. *Let $(\sigma^k)_{k=0}^{\hat{K}}$ be a behavior LPS. Then, the ℓ-strategy profile $\lambda(a) \equiv (\hat{K}(a), \sigma^{\hat{K}(a)}(a))$ for all $a \in A$ generates the same assessment.*

Proof. We first show that λ is an ℓ-strategy profile. To see this note that, for each information set I, if $a \in A^I$ is in the support of $\sigma^0|_{A^I}$, then $\hat{K}(a) = 0$, and if $a \in A^I$ is not in the support of $\sigma^0|_{A^I}$, then $\hat{K}(a) > 0$. It then follows that $\sum_{a\in A^I} \lambda(a) = (0, \sum_{a\in A^I} \sigma^0(a)) = 1$.

We proceed by showing that the behavior LPS and λ have the same outcome. It is clear that they generate the same strategy profile, since $\sigma(a) = \sigma^0(a) = \mathrm{st}(\lambda(a)) = \sigma^\lambda(a)$. Take then some information set I and history $h \in I$. Recall that

$$\lambda((a_j)_{j=1}^J) = \prod_{j=1}^J \lambda(a_j) = \left(\sum_{j=1}^J \hat{K}(a), \prod_{j=1}^J \sigma^{\hat{K}(a_j)}(a_j)\right) ,$$

and also that

$$\lambda(I) = \sum_{h\in I} \lambda(h) = \left(\lambda^{\mathrm{r}}(I), \sum_{h\in I|\lambda^{\mathrm{r}}(h)=\lambda^{\mathrm{r}}(I)} \lambda^{\mathrm{p}}(h)\right) .$$

Since $\lambda^{\mathrm{r}}(h) = \hat{K}(h)$ and $\lambda^{\mathrm{r}}(I) = \hat{K}(I)$ for all $h \in H$ and $I \in \mathcal{I}$, it is clear that $\mu(h) = 0$ if and only if $\mu^\lambda(h) = 0$. If, alternatively, $\hat{K}(h) = \hat{K}(I)$ (and so $\lambda^{\mathrm{r}}(h) = \lambda^{\mathrm{r}}(I)$), then

$$\mu((a)_{j=1}^J) = \frac{\prod_{j=1}^J \sigma^{\hat{K}(a_j)}}{\sum_{h\in I|\lambda^{\mathrm{r}}(h)=\lambda^{\mathrm{r}}(I)} \lambda^{\mathrm{p}}(h)} = \mathrm{st}\left(\frac{\lambda((a)_{j=1}^J)}{\lambda(I)}\right) = \mu^\lambda((a)_{j=1}^J) ,$$

hence the result holds. □

Proposition B.1 is illustrative of how an ℓ-strategy profile retains the information necessary to determine the likelihood of each action, which is given by its likelihood in the most likely theory where it is played with positive probability. This information is sufficient to determine the consistency and sequential optimality of assessments, as it permits assessing the relative likelihood of any two histories.

In applications, especially in extensive-form games, using LPSs to study behavior may require high values for $\hat{K}$, diminishing their usefulness. Still, as we argued, most of the components of an LPS are not necessary to determine the incentives of the players in the game. As a result, the reduced dimensionality of ℓ-strategy profiles ($\mathbb{R}^{2\,|A|}$ instead of $\mathbb{R}^{(\hat{K}+1)\,|A|}$ in a behavior LPS), their additive and multiplicative properties, and the fact that they can be defined in each information set independently, make them easier to work with in applications.[18]

B.2 Conditional Probability Systems (CPS)

A different approach to model consistent behavior is through the use of conditional probability systems, like Battigalli (1996), or relative probabilities, like Kohlberg and

[18] Note that Govindan and Klumpp (2003) use LPSs to characterize perfect equilibria. Nevertheless, their analysis requires using *induced lexicographic beliefs* for each player (which plays a similar role as the "independence property" for CPSs), which add a significant degree of intractability as they are $(\hat{K}\,(|N|-1))$-dimensional objects. Nevertheless, as we indicate in footnote 3, the equilibrium paths of perfect and sequential equilibria coincidence for generic payoffs.

Reny (1997). In this section we focus on the relationship of between conditional probability systems and ℓ-strategy profiles. To simplify the exposition, we assume that nature does not take any action.

Let $S \equiv \{s \in A^{\mathcal{I}} | s(I) \in A^I \ \forall I \in \mathcal{I}\}$ be the set of pure strategies. Battigalli (1996) defines *conditional probability system (CPS)* on the set of pure strategies S to be a map $P : 2^S \times (2^S \backslash \emptyset) \to [0,1]$ such that for all $S_1 \in 2^S \backslash \emptyset$ we have that $P(\cdot|S_1) \in \Delta(S_1)$, and for all $S_1, S_2, S_3 \subset S$,

$$S_1 \subset S_2 \subset S_3 \text{ implies } P(S_1|S_3) = P(S_1|S_2)\, P(S_2|S_3). \tag{B.2}$$

To ensure that a CPS is consistent with independent randomizations, Battigalli defines the following property: a CPS P *satisfies the independence property* if, for any partition $\{\mathcal{I}', \mathcal{I}''\}$ of $\mathcal{I}$ and any two sets $S_1' \times S_1'', S_2' \times S_2'' \in A^{\mathcal{I}'} \times A^{\mathcal{I}''}$, we have

$$P(S_1' \times S_1'' | S_2' \times S_1'') = P(S_1' \times S_2'' | S_2' \times S_2'') \ .$$

Our definition is slightly different than Battigalli's in that we define the independence property over behavior pure strategies instead of pure strategies (that is, we allow different information sets of the same player to be in different elements of the partition $\{\mathcal{I}', \mathcal{I}''\}$ of $\mathcal{I}$). In our context of an extensive form game with a focus on behavior strategies, this is without loss of generality, since each player randomizes independently in each of her information sets. Finally, Battigalli defines a *strategic extended assessment* as a triple (σ, μ, P), where (σ, μ) is an assessment and P is a CPS satisfying

$$\mu^I(h) = P\big(S(h) \big| S(I)\big) \quad \text{and} \quad \sigma(a) = P\big(\, \{s \in S(I^a) \mid s(I^a) = a\, \} \,\big|\, S(I^a)\,\big) \tag{B.3}$$

for all $I \in \mathcal{I}$, $h \in I$ and $a \in A$, where $S(h \equiv (a_j)_{j=1}^J)$ is the set of elements $s \in S$ such that $s(I^{a_j}) = a_j$ for all j (that is, h is on path of s), and $S(I) \equiv \cup_{h' \in I} S(h')$.

Now fix some ℓ-strategy profile λ. For any element $s \in S$ we define, with some abuse of notation, its likelihood as $\lambda(s) \equiv \prod_{I \in \mathcal{I}} \lambda(s(I))$. Also, for a given set $S_1 \subset S$, we use $\lambda(S_1)$ to denote $\sum_{s \in S_1} \lambda(s)$. Finally, the *conditional probability system generated by* λ, denoted P^λ, as

$$P^\lambda(S_1|S_2) \equiv \text{st}\left(\frac{\lambda(S_1 \cap S_2)}{\lambda(S_2)} \right) \quad \text{for all } S_1, S_2 \subset S \text{ with } S_2 \neq \emptyset \ . \tag{B.4}$$

The following result illustrates the connection of ℓ-strategy profilesand CPS:

Proposition B.2. *For each ℓ-strategy profile λ, P^λ satisfies the "independence property" and $(P^\lambda, \mu^\lambda, \sigma^\lambda)$ is a strategic extended assessment.*

Proof. It is clear that P^λ satisfies (B.2). To see that P^λ satisfies the independence property, consider a partition $\{\mathcal{I}', \mathcal{I}''\}$ of $\mathcal{I}$ and $S_1' \times S_1'', S_2' \times S_2'' \subset A^{\mathcal{I}'} \times A^{\mathcal{I}''}$. Then, note that

$$P^\lambda(S_1'\times S_1''|S_2'\times S_1'') = \frac{P^\lambda\big((S_1'\times S_1'')\cap(S_2'\times S_1'')\big)}{P^\lambda(S_2'\times S_1'')} = \frac{P^\lambda\big((S_1'\cap S_2')\times S_1''\big)}{P^\lambda(S_2'\times S_1'')}$$
$$= \frac{(\sum_{s'\in S_1'\cap S_2'}\lambda(s'))\,(\sum_{s''\in S_1''}\lambda(s''))}{(\sum_{s'\in S_1'\cap S_2'}\lambda(s'))\,(\sum_{s''\in S_1''}\lambda(s''))} = \frac{\lambda(S_1'\cap S_2')}{\lambda(S_2')}$$

does not depend on S_1'', where λ is naturally extended to $A^{\mathcal{I}'}$ and $A^{\mathcal{I}''}$.[19] Hence, P^λ satisfies the independence property. To prove that $(P^\lambda,\mu^\lambda,\sigma^\lambda)$ is a strategic extended assessment, notice that, for any history $h\equiv(a_j)_{j=1}^J$,

$$\sum_{s\in S(h)}\lambda(s) = \Big(\prod_{j=1}^J\lambda(a_j)\Big)\Big(\prod_{I\notin\{I^{a_j}\,|\,j=1,\dots,J\}}\underbrace{\sum_{a\in A^I}\lambda(a)}_{=1}\Big) = \lambda(h)\ . \tag{B.5}$$

Similarly, we have

$$\sum_{s\in S(I)}\lambda(s) = \sum_{h\in I}\sum_{s\in S(h)}\lambda(s) = \lambda(I)\ . \tag{B.6}$$

As a result, using the definitions of $(\mu^\lambda,\sigma^\lambda)$ (Definition 3.2) and P^λ (equation (B.4)), the first condition in equation (B.3) holds. To prove that the second condition in equation (B.3) holds, fix an action $a\in A$. Using $S(a)$ to denote $\{s\in S(I^a)\,|\,s(I^a)=a\}$, and an argument similar to the one used to obtain equations (B.5) and (B.6), we have

$$\sum_{s\in S(a)}\lambda(s) = \lambda(a)\sum_{h\in I^a}\lambda(h)\ . \tag{B.7}$$

Then, the second condition in equation (B.3) holds, and the proof is done. □

Corollary 3.1 in Battigalli (1996) shows that an assessment (μ,σ) is consistent only if there is some CPS P satisfying the independence property such that (μ,σ,P) is an strategic extended assessment. Proposition B.2 provides us with one of such CPSs: it is the one generated by an ℓ-strategy profile that generates (μ,σ).

CPSs and relative probabilities are difficult to use in applications. One reason is their high dimensionality ($2^{2|S|}$, in most games much higher than the set of ℓ-trembles, which is $2\,|A|$). Furthermore, the requirements that a given CPS satisfies the independence property or that a given pair assessment-CPS is a strategic extended assessment may be difficult to verify, given the large number of equations they entail.

Remark B.1. Kohlberg and Reny (1997) use the assessment of an "outside observer" to characterize the relative probability of any $|N|$-tuple of (normal-form game) pure

[19]For example, $\lambda(s')\equiv\prod_{I\in\mathcal{I}'}\lambda(s(I))$ for each $s'\in A^{\mathcal{I}'}$, and $\lambda(S')\equiv\sum_{s'\in S'}\lambda(s')$ for each $S'\subset A^{\mathcal{I}'}$.

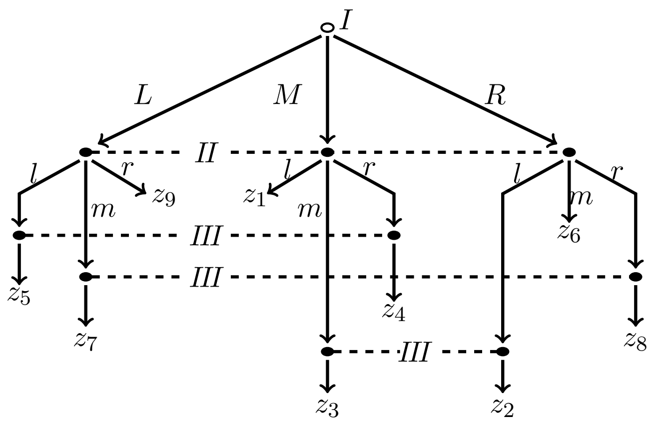

Figure 2: Tree from Example B.1.

strategies of the players. Our construction can also be interpreted as generating, for each pair of histories of the extensive form game, an outside observer's assessment about their relative likelihood. In fact, Proposition 3.1 shows that it is indeed the case: if $(\mu_n, \sigma_n)_n$ is a sequence of fully-mixed assessments converging to some assessment (μ, σ) if and only if there is an ℓ-strategy profile λ that generates (μ, σ), and satisfies[20]

$$\forall h, h' \in H, \quad \lim_{n\to\infty} \frac{\Pr(h|\sigma_n)}{\Pr(h'|\sigma_n)} = \begin{cases} 0 & \text{if } \lambda^{\mathrm{r}}(h) > \lambda^{\mathrm{r}}(h'), \\ \lambda^{\mathrm{p}}(h)/\lambda^{\mathrm{p}}(h') & \text{if } \lambda^{\mathrm{r}}(h) = \lambda^{\mathrm{r}}(h'), \\ \infty & \text{if } \lambda^{\mathrm{r}}(h) < \lambda^{\mathrm{r}}(h'). \end{cases} \tag{B.8}$$

Example B.1. Consider Figure 2, which is and example of how the independence property is not sufficient to guarantee consistency. Consider a CPS characterized by $P(z_n|z_{n'}) = 0$ for all $n > n'$. We illustrate how the use of ℓ-strategy profiles simplifies showing that, while P satisfies the independence property, it is not part of a sequential equilibrium.[21] Indeed, if P was part of a sequential equilibrium, there would exist some ℓ-strategy profile λ generating it. Then, we would have $\lambda^{\mathrm{r}}(M) = \lambda^{\mathrm{r}}(l) = 0$ (because $\lambda^{\mathrm{r}}(z_1) = 0$). Also, since $\lambda^{\mathrm{r}}(z_2) < \lambda^{\mathrm{r}}(z_3)$ and $\lambda^{\mathrm{r}}(z_4) < \lambda^{\mathrm{r}}(z_5)$, we have $\lambda^{\mathrm{r}}(R) < \lambda^{\mathrm{r}}(m)$ and $\lambda^{\mathrm{r}}(r) < \lambda^{\mathrm{r}}(L)$. Finally, since $\lambda^{\mathrm{r}}(z_7) < \lambda^{\mathrm{r}}(z_8)$, we have that $\lambda^{\mathrm{r}}(L) + \lambda^{\mathrm{r}}(m) < \lambda^{\mathrm{r}}(R) + \lambda^{\mathrm{r}}(r)$, which is a clear contradiction.

[20] In the same way that different sequences supporting a given consistent assessment give different limit relative probabilities between histories, different ℓ-strategy profilesgenerating the same assessment may also differ on the relative probability they assign to different histories. Still, in all of them, (B.8) holds whenever h and h' belong to the same information set.

[21] Our figure corresponds to Figure 2 in Battigalli (1996) and to Figure 7 in Kohlberg and Reny (1997). In both cases, showing that P is not part of a sequential equilibrium is involved because requires using sequences of positive (product) probabilities.