Lexicographic Numbers
in Extensive Form Games

Francesc Dilmé

Summer 2022

Abstract

This paper introduces the space of (lexicographic) ℓ-numbers and uses it to analyze extensive-form games. ℓ-Numbers are a simplified, two-dimensional version of the hyperreal numbers, with straightforward operations that make it easy to work with infinitesimal probabilities. We use ℓ-numbers to provide simple characterizations of tremble-based equilibrium concepts, such as sequential equilibria or sequential stable outcomes, without using sequences of strategy profiles.

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

JEL classification codes: C72, C73.

University of Bonn. fdilme@uni-bonn.de. I thank Sarah Auster, Stephan Lauermann, Srihari Govindan, Georg Nöldeke, and the seminar participants at the University of Bonn, Tel Aviv University, Hebrew University at Jerusalem, EEA-ESEM 2022 in Milan, for their useful comments. This work was funded by a grant from the European Research Council (ERC 949465), by the Deutsche Forschungsgemeinschaft (German Research Foundation) through grant CRC TR 224 (Project B02), and under Germany's Excellence Strategy (EXC 2126/1-390838866 and EXC 2047-390685813).

1 Introduction

Sequential rationality is a plausible requirement in extensive-form games. Broadly speaking, it requires players to behave optimally at any instance of the game. For example, in a standard signaling game, the sender should expect a rational receiver to best-respond to each message according to some beliefs, whether on path or not. If the receiver assumes that the sender is playing a given strategy, rationality uniquely pins down her assessment of the relative likelihood of each type of the sender after each on-path message; hence her best responses are determined as well. If the message is off path, rationality alone does not restrict the receiver's assessment. Such indefiniteness often generates a multiplicity of equilibria, some of them sustained through implausible threats.

An important class of equilibrium concepts addressing the indefiniteness of off-path beliefs is that of tremble-based equilibria. These are typically expressed as limits of equilibria of versions of the game where strategies are perturbed. For example, trembling-hand perfect equilibria (Selten, 1975) are defined as limits of equilibria of close-by games where players tremble and hence Bayes' rule can be used after all histories. The advantage is that, in each equilibrium of a perturbation of the game, beliefs are updated rationally and behavior is sequentially rational, giving plausibility and a sense of robustness to the limit behavior. Other prominent examples of tremble-based equilibrium concepts are proper equilibria (Myerson, 1978), sequential equilibria (Kreps and Wilson, 1982), strategically stable sets of equilibria (Kohlberg and Mertens, 1986), and sequential stable outcomes (Dilmé, 2022b).

In practice, tremble-based equilibrium concepts are difficult to use. The reason is that proving or disproving that a given candidate is indeed an equilibrium often requires manipulating sequences of perturbations and limits of corresponding equilibrium strategies, which is often not feasible in applications. Instead, weaker and less theoretically appealing selection criteria are often used, such as different versions of perfect Bayesian equilibrium, iterated deletion of weakly dominated strategies, or criteria specific to signaling games (such as the Intuitive Criterion, D1, divinity, etc.).

In this paper, we develop a new tool aimed at making tremble-based equilibrium concepts easier to use. We construct the space of *ℓ-numbers* in order to represent and easily work with asymptotic likelihoods without using sequences. We use ℓ-numbers to provide a simple characterization of belief consistency and sequential equilibria. We then study the natural extension of previous stability concepts using ℓ-numbers: We introduce *ℓ-stable outcomes* and prove they always exist and have desirable properties. We show that ℓ-stable outcomes are typically easier to find than outcomes satisfying other stability concepts, and that finding ℓ-stable outcomes permits finding sequential stable and stable outcomes in some cases.

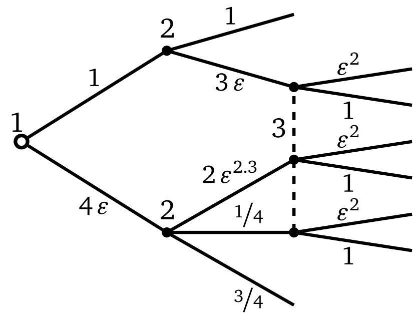

Figure 1

We begin by defining the space of (lexicographic) *ℓ-numbers*, also referred to as *likelihoods*, as the set of products of strictly positive real numbers and non-negative powers of an infinitesimal ε; that is, an ℓ-number is $x\,\varepsilon^y$ with $x \in \mathbb{R}_{++}$ and $y \in \mathbb{R}_+$. We identify ε^0 with 1, so ℓ-numbers are an extension of the positive real numbers. An ℓ-number $x\,\varepsilon^y$ with $y > 0$ is referred to as an *infinitesimal*, and it is used to categorize the likelihood of events with vanishing probability. Importantly, when adding ℓ-numbers we keep only the largest terms; for example, $1+4\,\varepsilon = 1$ and $5\,\varepsilon^{2.1}+2\,\varepsilon^{2.1} = 7\,\varepsilon^{2.1}$.[1] The standard part of an ℓ-number is its closest real number; for example, $\mathrm{st}(3\,\varepsilon^2) = 0$ and $\mathrm{st}(0.7\,\varepsilon^0) = 0.7$.

We use the language of ℓ-numbers to generalize the concept of a 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. For example, consider the game tree in Figure 1 and the corresponding ℓ-strategy profile (note that the sum of likelihoods of the actions at each node is 1). The likelihood of the middle history of player 3's information set is computed multiplying the likelihoods of the actions it contains, that is, $4\,\varepsilon \cdot 2\,\varepsilon^{2.3} = 8\,\varepsilon^{3.3}$. The likelihood of player 3's information set is computed by adding the likelihoods of its histories, that is, $3\,\varepsilon + 8\,\varepsilon^{3.3} + \varepsilon = 4\,\varepsilon$. The conditional likelihood of the middle history is then $8\,\varepsilon^{3.3}/(4\,\varepsilon) = 2\,\varepsilon^{2.3}$. Note that the ℓ-strategy profile generates a unique assessment (i.e., strategy profile and belief system), where each action is played with probability equal to the standard part of the assigned likelihood, and each history is assigned the standard part of its conditional likelihood. For example, under the assessment corresponding to Figure 1, player 1 plays her upper and lower actions with probability 1 and 0, respectively, and player 3's beliefs over the upper, middle, and lower histories are 0.75, 0, and 0.25, respectively.

[1] Because of this simplification, our construction can be interpreted as a significantly simplified version of hyperreal numbers (see Keisler, 2013, Henle and Kleinberg, 2014, and Robinson, 2016).

Our first result establishes that, despite their simplicity, ℓ-numbers permit a simple, full characterization of consistency (recall Kreps and Wilson, 1982, define consistent assessments as the limits of assessments of full-support strategy profiles). We show that an assessment is consistent if and only if some ℓ-strategy profile generates it. This result eases showing that a given assessment is consistent (it only requires providing ℓ-strategy profile generating it) or ruling out its consistency (see Example 3.4 below). Furthermore, it allows us to characterize sequential equilibria. We define an *ℓ-equilibrium* as an ℓ-strategy profile where the likelihood of an action is not infinitesimal only if the action is sequentially optimal given the implied belief system. We show that each ℓ-equilibrium generates a sequential equilibrium and that each sequential equilibrium is generated by some ℓ-equilibrium. This characterization makes it easier to work with sequential equilibria, as it permits one to work directly "in the limit" to determine *both* sequential optimality and belief consistency.

With the goal of studying stability in mind, we proceed by generalizing the concept of a tremble sequence in Section 4. An *ℓ-tremble* assigns an infinitesimal to each action and is interpreted as a very small likelihood with which players make mistakes. 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 sequentially optimal actions. We provide a characterization of ℓ-equilibria for a given ℓ-tremble as those which generate assessments that are limits of perfect ε_n-equilibria of versions of the game perturbed according to the ℓ-tremble, for some sequence $\varepsilon_n \to 0$.

We then define stability using ℓ-numbers by requiring stability with respect to ℓ-trembles: 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. We show that sequential stable outcomes (Dilmé, 2022b) are ℓ-stable, and hence all games have an ℓ-stable outcome. Like sequential stable outcomes, ℓ-stable outcomes satisfy the never a weak best response (NWBR, which implies versions forward induction and iterated strict dominance), invariance to reordering simultaneous moves, and on-path subgame perfection. They also pass standard selection criteria in signaling games (e.g., Intuitive Criterion, D1, and D2).

We finally provide procedures for obtaining ℓ-stable outcomes in practice. The first consists in ruling out the ℓ-stability of a given outcome ω by (i) proving that it fails some necessary condition (e.g., NWBR, sequential rationality, or D1), or (ii) showing that there is no ℓ-equilibrium for a given ℓ-tremble with outcome ω, or (iii) a combination of the two. If all but one outcome are shown to *not* be ℓ-stable, then existence implies that the remaining outcome is ℓ-stable. The second procedure consists in explicitly showing that any ℓ-tremble has an ℓ-equilibrium with a given outcome. Both procedures are greatly simplified when using ℓ-numbers instead of sequences of trembles and corresponding sequences of strategy profiles. We show that ℓ-strategy profiles can be generalized to allow for 0-likelihood actions, which simplifies further working with ℓ-numbers. We illustrate these procedures through some examples.

1.1 Related literature

Our paper contributes to the literature aiming to simplify the use of tremble-based equilibrium concepts. In our view, there are three main approaches, which we now describe.

The first approach consists in analyzing weaker equilibrium concepts that are easier to use. For example, Kreps and Wilson (1982) introduced sequential equilibria as a weaker but easier-to-use concept than the perfect equilibria of Selten (1975): The former required beliefs to be consistent with a sequence of strategy profiles, but sequential optimality was required only "in the limit" instead of along the sequence. (As Kreps and Wilson state, "It is vastly easier to verify that a given equilibrium is sequential than that it is perfect.") Perfect Bayesian equilibrium (Fudenberg and Tirole, 1991) was a further weakening, with a weaker notion of belief consistency that did not require using sequences of strategy profiles. Similarly, Dilmé (2022b) weakened the concept of stable outcomes in Kohlberg and Mertens (1986) to the easier-to-use concept of sequential stable outcome by requiring ε-optimality along the sequence instead of exact optimality.

The second approach consists in using selection criteria, which are "tests" that permit one to eliminate undesirable equilibria. This approach has been popular in signaling games, where selection criteria (such as the Intuitive Criterion, D1, and D2 of Cho and Kreps, 1987, and the divinity criterion of Banks and Sobel, 1987) were introduced to simplify obtaining stable outcomes. Similarly, equilibria are sometimes ruled out because they fail to satisfy certain properties of tremble-based equilibria (such as forward induction or "no signaling what you do not know").

The last approach consists in generalizing the definition of strategy profiles or belief systems to richer mathematical objects. Myerson (1986) generalizes belief systems to conditional probability systems (CPSs) to specify conditional probabilities (or beliefs) on zero-probability events. He shows that CPSs are limits of sequences of probability distributions and uses them to characterize sequential communication equilibria and predominant communication equilibria (see also Battigalli, 1996, and Kohlberg and Reny, 1997). Blume, Brandenburger, and Dekel (1991) instead generalize strategy profiles to lexicographic probability systems (LPSs), provide a decision-theoretic representation of preferences under lexicographic beliefs, and use it to characterize (normal-form) perfect and proper equilibria (see also Govindan and Klumpp, 2003). Similarly, Mailath, Samuelson, and Swinkels (1997) use LPSs to characterize strategic independence-respecting equilibria (SIRE) and compare them to proper equilibria.

This paper adopts the third approach. We deliberately construct the simplest generalization of strategy profiles that permits the characterization of tremble-based equilibrium concepts without using sequences. Our likelihoods are two-dimensional, assigned separately to each action, and can be manipulated using simple operations. While the construction permits us to characterize weaker concepts than LPSs—sequential instead of perfect/proper equilibria—it is also significantly easier to use, which makes it useful for characterizing (sequential) stability. Appendix B contains a detailed exposition of the relationship between ℓ-strategy profiles, LPSs, and CPSs.

The rest of the paper is organized as follows. Section 2 introduces the space of ℓ-numbers. Section 3 defines ℓ-strategy profiles and ℓ-equilibria. It also characterizes consistent assessments using ℓ-strategy profiles and sequential equilibria using ℓ-equilibria. Section 4 defines ℓ-stable outcomes, proves their existence, and describes some of their properties. It also provides some results and techniques to find ℓ-stable outcomes, relates them to common selection criteria in signaling games, and illustrates how to use them through examples. Finally, Section 5 concludes. The appendix contains the proofs of the results and discusses the relationship between ℓ-strategy profiles, LPSs, and CPSs.

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 will do so by using a simplified version of the usual calculus with infinitesimals (see Keisler, 2013, Henle and Kleinberg, 2014, and Robinson, 2016).

From now on, the symbol ε will signify an "infinitesimal," which informally can be thought of as a number bigger than 0 but smaller than any positive real number. An *ℓ-number* ℓ is the product of a positive real number $x \in \mathbb{R}_{++}$ and ε to the power of a non-negative real number $y \in \mathbb{R}_{+}$, that is, $\ell \equiv x\,\varepsilon^{y}$. The space of ℓ-numbers is L. We will use ℓ-numbers to represent asymptotic likelihoods with which actions are played along some sequence of strategy profiles. We now define simple operations for ℓ-numbers.

Definition 2.1. We endow the space of ℓ-numbers with the following order and operations for each pair $\ell \equiv x\,\varepsilon^y, \hat{\ell} \equiv \hat{x}\,\varepsilon^{\hat{y}} \in L$:

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

We identify each ℓ-number $x\,\varepsilon^0 \in L$ with the real number $x \in \mathbb{R}_{++}$ (i.e., $\varepsilon^0 \equiv 1$), so the previous operations are consistent with the usual addition, multiplication, and division of real numbers. ℓ-Numbers are therefore an extension of the positive numbers, $\mathbb{R}_{++} \subset L$. Also, we call an ℓ-number $\ell \in L$ with $\mathrm{st}(\ell) = 0$ (i.e., with $y > 0$) an *infinitesimal*. We adopt the natural convention that $\varepsilon^1 \equiv \varepsilon$.

To gain intuition about the operations between ℓ-numbers, fix two ℓ-numbers $\ell, \hat{\ell} \in L$. The lexicographic order uses that, when comparing two ℓ-numbers, the one with the smallest power to ε is the largest one—as it would happen if ε were a very small but strictly positive real number. The multiplying factors are compared if the two powers are the same. Importantly, the sum of ℓ and $\hat{\ell}$ only preserves the most likely terms. For example, if $y < \hat{y}$ then $\ell + \hat{\ell} = x\,\varepsilon^y$, and if $y = \hat{y}$ then $\ell + \hat{\ell} = (x + \hat{x})\,\varepsilon^y$. This assumption (which differs from the standard use of hyperreal numbers, see below) permits keeping the space ℓ-numbers simple while still allowing to compare the likelihood of low probability events. Multiplication and division of ℓ-numbers are also simple and intuitive, while the standard part simply provides the closest real number. As it will become apparent, it is convenient (and unnecessary) for our analysis not to include 0 in the set of ℓ-numbers. Still, as we argue in Section 4.4, including 0 has some tractability advantages in some games; in there, we present an extension of the space of ℓ-numbers that contains the number 0.

Representation as classes of power sequences

It will be sometimes convenient to think of an ℓ-number $\ell = x\,\varepsilon^y$ as indexing the (equivalence) class of sequences that tend to 0 as $(x\,(1/n)^y)_n$. That is, a given real sequence $(\sigma_n(a))_n$, where $\sigma_n(a) \in \mathbb{R}_{++}$ for all $n \in \mathbb{N}$, belongs to the class indexed by ℓ, denoted $(\sigma_n(a))_n \in \ell$, if and only if

$$\lim_{n\to\infty} \frac{\sigma_n(a)}{x\,(1/n)^y} = 1\ . \tag{2.1}$$

This representation is helpful to get further intuition for the operations defined in Definition 2.1. Take two ℓ-numbers $\ell, \hat{\ell} \in L$, and let $(\sigma_n(a))_n \in \ell$ and $(\hat{\sigma}_n(a))_n \in \hat{\ell}$ be two sequences. Then, we have

$$\ell > \hat{\ell} \quad \Leftrightarrow \quad \lim_{n\to\infty} \frac{\sigma_n(a)}{\hat{\sigma}_n(a)} > 1 \ .$$

That is, $\ell > \hat{\ell}$ if any sequence represented by ℓ is asymptotically bigger than any sequence represented by $\hat{\ell}$. Similarly, the operations defined above correspond to operations between sequences they represent; for each of the operations $\star \in \{+, \cdot, /\}$, we have

$$(\sigma_n(a) \star \hat{\sigma}_n(a))_n \in \ell \star \hat{\ell} \ .$$

Intuitively, our definition of sum only preserves the most likely terms of $(\sigma_n(a) + \hat{\sigma}_n(a))_n$: if $y < \hat{y}$ then $(\sigma_n(a))_n$ tends to 0 at a slower rate than $(\hat{\sigma}_n(a))_n$, so $(\sigma_n(a) + \hat{\sigma}_n(a))_n$ tends to 0 at the same rate as $(\sigma_n(a))_n$, hence $\ell + \hat{\ell} = \ell$. The standard part of an ℓ-number ℓ is the limit of the sequences it contains, $\mathrm{st}(\ell) = \lim_{n\to\infty} \sigma_n(a)$: $\mathrm{st}(\ell) = 0$ if $y > 0$ and $\mathrm{st}(\ell) = x$ if $y = 0$.[2]

Relationship to hyperreal numbers

There is abundant work in mathematics on hyperreal numbers and their use to study calculus using infinitesimals. The set of *hyperreal numbers* is an extension of the real numbers, constructed through the addition of "infinitesimals" and "infinites" and the requirement of closeness through basic operations and limits (see Keisler, 2013, for a formal definition).

Working with hyperreal numbers is convenient to simplify some limit arguments in calculus, such as approximations of functions around a given point. Hyperreal numbers are nevertheless difficult to use in practice because of their infinite dimensionality, as multiple infinitesimals are kept through the operations and limits. For example, the sum and multiplication of the hyperreal numbers $1 + \varepsilon$ and $e^{-1/\varepsilon} - 1/\log(\varepsilon)$ are

$$1 + \varepsilon + e^{-1/\varepsilon} - 1/\log(\varepsilon) \ \text{ and } \ e^{-1/\varepsilon} - 1/\log(\varepsilon) + \varepsilon\, e^{-1/\varepsilon} - \varepsilon/\log(\varepsilon) \ ,$$

respectively. Then, through the sum and multiplication (and other operations), the representation of hyperreal numbers gets typically long and unusable in many settings.

Our construction of ℓ-numbers is tailored to be used to model probabilities in the study of game theory and greatly simplifies hyperreal numbers. The first simplification is that we only consider

[2] Note that the representation can be generalized by replacing $x\,(1/n)^y$ by $x\,f(n)^y$ in expression (2.1), where $f : \mathbb{N} \to \mathbb{R}_{++}$ is any function such that $\lim_{n\to\infty} f(n) = 0$. We take $f(n) = 1/n$ for concreteness and simplicity.

infinitesimals that are (positive) powers of some fixed infinitesimal ε. Second, we disregard small infinitesimals when we use addition, hence keeping the expressions simple. As we shall now see, despite its simplicity, our construction will be powerful in characterizing sequential equilibria and stable outcomes.

3 ℓ-Strategy profiles and belief consistency

In this section, we use ℓ-numbers to characterize belief consistency and sequential optimality in extensive-form games. We first establish the notation for extensive-form games. We then introduce ℓ-strategy profiles and prove that the set of assessments they generate coincides with the set of consistent assessments. We finally introduce the concept of ℓ-equilibrium and show that it is equivalent to the concept of sequential equilibrium.

3.1 Extensive-form games

We now provide the definitions and corresponding notation for an extensive-form game.

A (finite) *extensive-form game* $G=\langle A,H,\mathcal{I},N,\pi,u\rangle$ has the following components. A finite set of *actions* A. A finite set of *histories* H containing finite sequences of actions, satisfying that for all $(a_j)_{j=1}^J\in H$ with $J>0$ we have $(a_j)_{j=1}^{J-1}\in H$; the set of *terminal histories* is denoted T. An *information partition* $\mathcal{I}$ of the non-terminal histories and a partition $\{A^I|I\in\mathcal{I}\}$ of A such that, for each $I\in\mathcal{I}$ and $h\in H$, we have (i) $(h,a)\in H$ for some $a\in A^I$ if and only if $h\in I$, and (ii) if $h\in I$ and $h'>h$ then $h'\notin I$.[3] A finite set of *players* $N\not\ni 0$. A *player assignment* $\iota:\mathcal{I}\to N\cup\{0\}$ assigning each information set to a player or to nature. A *strategy by nature* $\pi:\cup_{I\in\iota^{-1}(\{0\})}A^I\to(0,1]$ satisfying $\sum_{a\in A^I}\pi(a)=1$ for each $I\in\iota^{-1}(\{0\})$. For each player $i\in N$, a *(von Neumann–Morgenstern) payoff function* $u_i:T\to\mathbb{R}$. For convenience, we set $u_0(t)=0$ for all $t\in T$.

A *strategy profile* is a map $\sigma:A\to[0,1]$ such that $\sum_{a\in A^I}\sigma(a)=1$ for all $I\in\mathcal{I}$ (i.e., it is a probability distribution for each set of actions available at each information set) and $\sigma(a)=\pi(a)$ for all a played by nature (i.e., nature plays according to π). We let Σ be the set of strategy profiles.

[3] Note that 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.

3.2 ℓ-Strategy profiles

We begin by defining ℓ-strategy profiles. 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 $\lambda : A \to L$ satisfying that, for each $I \in \mathcal{I}$, (i) $\sum_{a \in A^I} \lambda(a) = 1$, and (ii) $\lambda(a) = \pi(a)$ for all $a \in A^I$ whenever $\iota(I) = 0$.

The set of ℓ-strategy profiles is denoted Λ. An ℓ-strategy profile satisfies the same conditions as a strategy profile. The first condition requires that the total likelihood of the actions of a given information set is 1. The second condition is that actions by nature receive the same likelihood as the one specified by the game. Without risk of confusion, we will often use $x(a)\,\varepsilon^{y(a)}$ to denote $\lambda(a)$.

We can easily define the likelihood of a history or an information set induced by an ℓ-strategy profile $\lambda \in \Lambda$ 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)

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

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.[4]

We do not interpret an ℓ-strategy profile λ as a profile of choices of sequences of strategy profiles by the players. Instead, in the spirit of Kohlberg and Reny (1997)'s interpretation of relative probability spaces, each ℓ-number $\lambda(a)$ can be seen as a simple tool used by an outside observer to assess the relative likelihoods of actions, histories, and information sets. The outside observer can then "consistently" use the ℓ-strategy profile to assess the relative likelihood of histories. As we shall now see, all and only assessments generated from ℓ-strategy profiles will satisfy Kreps and Wilson (1982)'s consistency condition.

Remark 3.1. Our definitions of ℓ-numbers and ℓ-strategy profiles do not allow actions to be played with likelihood (exactly) 0. In Section 4.1 we justify this restriction by introducing ℓ-trembles, which set a lower bound on the likelihood each action can be played. There are, nevertheless, other possible definitions of ℓ-numbers and ℓ-strategy profiles which allow the likelihood of an action to be

[4]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 $y(h) < y(h')$ (so h is infinitely more likely than h'), then $\lambda(I) = \lambda(h) + \lambda(h') = \lambda(h)$.

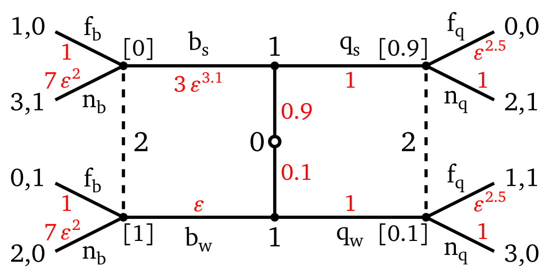

Figure 2

exactly 0 while preserving all our results. In Section 4.4, we show that all our results hold under the less restrictive definition of *(fully-mixed) ℓ-strategy profile* (see Definition 4.5), which allows $\lambda(a)=0$ for some actions $a\in A$ while ensuring that information sets are reached with non-zero (but maybe infinitesimal) likelihood. We argue that such relaxation simplifies some arguments.

Example 3.1. Figure 2 depicts the beer-quiche game (where nature plays according to $\pi(\mathrm{s})=1-\pi(\mathrm{w})=0.9$) and an ℓ-strategy profile λ (in red). Using I_b to denote player 2's left information set, we have

$$\lambda(I_\mathrm{b}) = \lambda(\mathrm{s},\mathrm{b_s})+\lambda(\mathrm{w},\mathrm{b_w}) = 2.7\,\varepsilon^{3.1}+0.1\,\varepsilon = 0.1\,\varepsilon\ .$$

3.3 Assessments and consistency

We continue by relating ℓ-strategy profiles 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 strategy profile, and where $\mu : H\backslash T\to[0,1]$ is a belief system satisfying that $\sum_{h\in I}\mu(h)=1$ for all $I\in\mathcal{I}$. We can associate an assessment to an ℓ-strategy profile by taking the standard part of the likelihoods of each action and the relative likelihoods of histories as follows:

Definition 3.2. The assessment *generated* by $\lambda\in\Lambda$ is $(\sigma^\lambda,\mu^\lambda)$ 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 information set I, it is natural to require Bayes consistency if I is on path, and also to require that beliefs are updated consistently off path. Kreps and Wilson say that 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 to compute limits of strategy profiles.

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

The "if" part of Proposition 3.1 is obtained as follows. Fix some ℓ-strategy profile $\lambda \in \Lambda$ and define $(\sigma_n^\lambda)_n$ using power sequences of strategy profiles as follows: for each $n \in \mathbb{N}$ and $a \in A$,

$$\sigma_n^\lambda(a) \equiv \begin{cases} x(a)\,(1/n)^{y(a)} & \text{if } y(a) > 0, \\ M_n^\lambda(I^a)\,x(a) & \text{if } y(a) = 0, \end{cases} \tag{3.1}$$

where $M_n^\lambda(I^a)$ is a factor that ensures that $\sum_{a' \in A^{I^a}} \sigma_n^\lambda(a') = 1$, where $I^a \in \mathcal{I}$ is the unique information set where a is available (i.e., satisfying $a \in A^{I^a}$), and where the subindex n is 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 easy to see that $(\sigma_n^\lambda)_n$ supports $(\sigma^\lambda, \mu^\lambda)$, hence $(\sigma^\lambda, \mu^\lambda)$ is consistent for all $\lambda \in \Lambda$.

The "only if" part is less obvious, as a consistent assessment can be generated by strategy profiles that do not necessarily look like those in equation (3.1). The result follows from Theorem 3.1 in Dilmé (2022a), which shows that any consistent assessment is supported by a power sequence of strategy profiles like (3.1). That is, even though power sequences of strategy profiles are small compared to the set of all sequences, they are rich enough to support all consistent assessments. ℓ-Strategy profiles constitute then a simple tool to characterize consistency: they are defined action-by-action and generate all and only consistent assessments. Examples 3.4 and 3.5 illustrate how ℓ-strategy profiles can also be used to rule out the consistency of some assessments.

Example 3.2 (Continuation of Example 3.1). Letting λ be the ℓ-strategy profile in Figure 2, we have that $\sigma^\lambda(\mathrm{s}) = 1 - \sigma^\lambda(\mathrm{w}) = 0.9$, $\sigma^\lambda(\mathrm{q_s}) = \sigma^\lambda(\mathrm{q_w}) = \sigma^\lambda(\mathrm{n_q}) = \sigma^\lambda(\mathrm{f_b}) = 1$, and $\sigma^\lambda(a) = 0$ for all other actions. Also, the implied belief system is depicted in the figure (numbers in brackets). For example,

$$\mu^\lambda(\mathrm{s}, \mathrm{b_s}) = \mathrm{st}(\lambda(\mathrm{s}, \mathrm{b_s})/\lambda(I_\mathrm{b})) = \mathrm{st}(2.7\,\varepsilon^{3.1}/(0.1\,\varepsilon)) = \mathrm{st}(27\,\varepsilon^{2.1}) = 0\ .$$

Remark 3.2. In our construction, the player acting at a given information set can be thought of as only entertaining histories with a non-infinitesimal conditional likelihood. Indeed, even when an information set I is off path (i.e., $y(I) > 0$), we have that $\sum_{h \in I} \lambda(h)/\lambda(I) = 1$, and also $\mu^\lambda(h) = 0$ if and only if $y(h) > y(I)$.

3.4 ℓ-Equilibria and sequential equilibria

Payoffs and ℓ-equilibria

We first 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

$$u(a|\lambda) \equiv \sum_{t \in T^a} \mathrm{st}\left(\frac{\lambda(t)}{\lambda(a)\,\lambda(I^a)} \right) u_{\iota(I^a)}(t) \tag{3.2}$$

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). It is easy to see that $u(a|\lambda)$ coincides with the continuation payoff player $\iota(I^a)$ obtains from playing a under the assessment $(\sigma^\lambda, \mu^\lambda)$, which we denote $u(a|\sigma^\lambda, \mu^\lambda)$.[5]

We can now define 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 only sequentially optimal actions are played with positive probability. Because ℓ-strategy profiles generate assessments, the sequential optimality requirement can be applied to all information sets.[6]

Relationship to sequential equilibria

We now relate the concept of ℓ-equilibrium to the well-known concept of sequential equilibrium by Kreps and Wilson (1982). They say an assessment (σ, μ) to be a *sequential equilibrium* if it is consistent (as defined above) and sequentially rational (so $\sigma(a) > 0$ only if $u(a|\sigma,\mu) \geq u(a'|\sigma,\mu)$ for all $a' \in I^a$). 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 some ℓ-equilibrium generates it. 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

[5] Recall that, for a given assessment (σ, μ) and each action $a \in A$, the continuation outcome is uniquely determined, and so is then the continuation payoff of the player playing a.

[6] To some degree, our definition is analogous to Myerson (1978)'s characterization of trembling perfect equilibrium as the limit of ε_n-perfect equilibria as $\varepsilon_n \to 0$, where σ_n is an ε_n-perfect equilibrium if it is fully mixed and satisfies that $\sigma_n(a) > \varepsilon_n$ only if a is optimal.

finding a sequence of fully-mixed assessments that converges to (σ,μ). Such difficulties have favored the use of equilibrium concepts that are not as powerful in selecting Nash equilibria, such as perfect Bayesian equilibria (PBE, Fudenberg and Tirole, 1991). ℓ-Equilibria provide a simple characterization of sequential equilibria: Any ℓ-strategy profile satisfying sequential rationality generates a sequential equilibrium.[7]

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: a sequential equilibrium exists (by Kreps and Wilson, 1982) and it is generated by some ℓ-strategy profile (by Proposition 3.1); hence such an ℓ-strategy profile is an ℓ-equilibrium.

Example 3.3 (Continuation of Example 3.2). It is easy to see that the ℓ-strategy profile depicted in Figure 2 is an ℓ-equilibrium, hence it generates a sequential equilibrium.

3.5 Using ℓ-strategy profiles

Powers vs. multiplying factors

The belief system μ^λ generated by an ℓ-strategy profile $\lambda\in\Lambda$ provides the relative likelihoods of histories $h,h'\in H$ in the same information set I, denoted $\lambda(h)\equiv x(h)\,\varepsilon^{y(h)}$ and $\lambda(h')\equiv x(h')\,\varepsilon^{y(h')}$, where note that $x(h)$ and $y(h)$ can be computed independently; for $h=(a_j)_{j=1}^J$ we have

$$x(h)=\prod_{j=1}^J x(a_j) \text{ and } y(h)=\sum_{j=1}^J y(a_j)\,.$$

Computing such relative likelihood from λ can be done in two steps. The first step compares the powers of ε, that is, $y(h)$ and $y(h')$. If $y(h)>y(h')$ then $\mu^\lambda(h)=0$, and if $y(h)<y(h')$ then $\mu^\lambda(h')=0$. The second step applies only when $y(h)=y(h')$, in which case the relative likelihood is determined by the ratio the multiplying factors, $x(h)/x(h')$.

The previous observations simplify proving or disproving that a given assessment is consistent and hence has the potential of being part of a sequential equilibrium, as illustrated in Examples 3.4, 3.5, and 4.4.[8] For instance, it is easy to see that an assessment (μ,σ) such that $\mu(h)\in\{0,1\}$ for all $h\in H\backslash T$ can be generated by an ℓ-strategy profile with $x(a)=1$ for all $a\in A$.

[7] Our characterization eases proving in applications the claim (often stated in footnotes without proof) that a given PBE (σ,μ) is also a sequential equilibrium, as it only requires reporting appropiate ℓ-numbers for the actions a with $\sigma(a)=0$ that generate belief system μ.

[8] Dilmé (2022a) makes a similar observation when computing assessments derived from power sequences of strategy profiles.

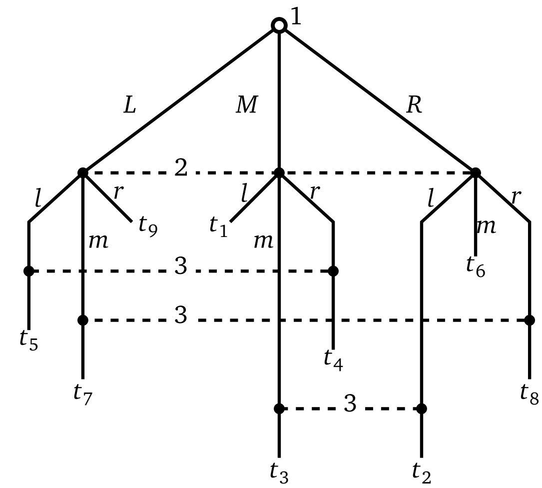

Figure 3

Example 3.4. Consider Figure 3, which corresponds to Figure 2 in Battigalli (1996) and to Figure 7 in Kohlberg and Reny (1997), with slight notation changes (see the discussion of conditional probability systems in Appendix B). Battigalli goes through some lengthy arguments to show that there is no sequence of strategy profiles $(\sigma_n)_n$ such that $\lim_{n\to\infty} \sigma_n(t_i)/\sigma_n(t_{i'})=0$ for all $i>i'$. Showing this is easy using Theorem 3.1: If $(\sigma_n)_n$ with the previous property existed, then there would be some mapping $y: A\to\mathbb{R}_+$ such that $y(t_i)>y(t_{i'})$ for all $i>i'$. Then, we would have $y(M)=y(l)=0$ (because $y(t_1)=0$). Also, since $y(t_2)<y(t_3)$ and $y(t_4)<y(t_5)$, we have $y(R)<y(m)$ and $y(r)<y(L)$. Finally, since $y(t_7)<y(t_8)$, we have that $y(L)+y(m)<y(R)+y(r)$, which is a contradiction.

Example 3.5. Consider again Figure 3. We prove that there is no consistent assessment where $\mu(L,l)=\mu(M,m)=\mu(R,r)=2/3$ (Kohlberg and Reny, 1997, prove this using limits of sequences). If this was the case, then there would be some map $x: A\to\mathbb{R}_{++}$ satisfying

$$x(L)\,x(l) = 2\,x(M)\,x(l), \quad x(M)\,x(m) = 2\,x(R)\,x(l), \quad \text{and}$$
$$x(R)\,x(l) = 2\,x(L)\,x(m).$$

The multiplication of all left-hand sides divided by the multiplication of all right-hand sides equals 1/8, which is an obvious contradiction.

Transformations

Different ℓ-strategy profiles can be equivalent in that they generate the same assessment. For example, if $\lambda\in\Lambda$ is such that there a single action a with $\text{st}(\lambda(a))=0$, then all ℓ-strategy profiles that coincide with λ except for the infinitesimal they assign to a generate the same assessment. Since some ℓ-strategy profiles may be easier to use than others equivalent to them, we now provide a simple transformation that preserves the assessment generated by a ℓ-strategy profile.

A straightforward way to transform an ℓ-strategy profile into an equivalent ℓ-strategy profile is by replacing the infinitesimal ε by $k_1\,\varepsilon^{k_2}$. Using this, it is easy to see that λ is equivalent to any ℓ-strategy profile obtained by choosing some $k_1, k_2 \in \mathbb{R}_{++}$ and applying

$$x(a) \mapsto k_1^{y(a)}\,x(a) \;\text{ and }\; y(a) \mapsto k_2\,y(a)$$

for all a, is equivalent to λ. An implication is that if a consistent assessment (μ, σ) satisfies $\sigma(a)=0$ for some $a \in A$, then (μ, σ) is generated by an ℓ-strategy profile with $\lambda(a)=\varepsilon$.

4 ℓ-Stable outcomes

Our next goal is to study stability using ℓ-numbers. We first introduce the concepts of ℓ-tremble and ℓ-equilibrium for an ℓ-tremble in this section. We will then define ℓ-stable outcomes and characterize their main properties. We alsoprovide some procedures and tools that simplify obtaining ℓ-stable outcomes. We then show that ℓ-strategy profiles can be extended to allow actions to have a likelihood equal to 0, which eases some arguments. We also present a renormalization of ℓ-strategy profiles that preserves their properties. Section 4.5 illustrates the use of ℓ-stability through an example.

4.1 ℓ-Trembles and ℓ-stable outcomes

ℓ-Equilibria for a given ℓ-tremble

Selten (1975) introduced the possibility that players may make mistakes. The idea was that, in a game where players tremble so that any action may be chosen by mistake, all information sets are on path; hence players can use Bayes rule to compute their beliefs. We now use the language of ℓ-numbers to model small probabilities of mistakes.

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

The set of ℓ-trembles is $\tilde{\Lambda}$. We interpret an ℓ-tremble as an infinitesimal likelihood with which players make mistakes (see below a characterization in terms of sequences of trembles). Note that each action is assigned an infinitesimal, and hence the standard part of an ℓ-tremble is not a probability distribution when restricted to the actions available at a given information set. We will often use $\tilde{x}(a)\,\varepsilon^{\tilde{y}(a)}$ to denote $\tilde{\lambda}(a)$, where note that $\tilde{y}(a) > 0$ for all $a \in A$.

Selten defined (trembling-hand) perfect equilibria as limits of equilibria of perturbed games. A strategy profile σ is a perfect equilibrium if there is a sequence of trembles $(\eta_n)_n \to 0$ and a corresponding sequence of strategy profiles $(\sigma_n)_n$ such that $\sigma_n(a) > \eta_n(a)$ only if a is sequentially optimal. The idea is that, in each perturbed game, each player plays an action with a probability no lower than the trembling probability and plays an action with a higher probability only if the action is optimal. We now impose the exact same requirements in our definition of ℓ-equilibrium for a given ℓ-tremble.

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

$$\lambda(a) > \tilde{\lambda}(a) \quad \Rightarrow \quad u(a|\lambda) \geq u(a'|\lambda) \;\text{ for all } a' \in A^{I^a}\,. \tag{4.1}$$

We denote the set of ℓ-equilibria for a given $\tilde{\lambda}$ as $\Lambda^*(\tilde{\lambda})$. As mentioned above, $\lambda \in \Lambda^*(\tilde{\lambda})$ satisfies two requirements. The first is that $\lambda(a) \geq \tilde{\lambda}(a)$ for all a, which 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 $\tilde{\lambda}(a)$. The second requirement is that $\lambda(a) > \tilde{\lambda}(a)$ only if a is sequentially optimal. Like for perfect equilibria, a player plays an action with a higher likelihood only if the action is optimal.

Note that the definition of ℓ-equilibrium (Definition 3.3) only requires sequential optimality for actions played with non-infinitesimal likelihood and does not impose any restriction on the infinitesimal assigned to the other actions. As we saw, this definition makes ℓ-equilibria equivalent to sequential equilibria. The definition ℓ-equilibria for a given ℓ-tremble (Definition 4.2) requires sequential optimality for all actions played with a likelihood strictly higher than that under the ℓ-tremble, even when their likelihood is infinitesimal. The following result establishes that, while the set of ℓ-equilibria for a particular ℓ-tremble is smaller than the set of ℓ-equilibria, any ℓ-equilibrium is an ℓ-equilibrium for some ℓ-tremble.

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

Proposition 4.1 establishes that the class of ℓ-equilibria that are consistent with optimization given some ℓ-tremble coincides with the class of all ℓ-equilibria. 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 $\text{st}(\lambda(a)) = 0$ and $\tilde{\lambda}(a) \equiv \varepsilon$ otherwise. The "only if" part is implied by the fact that, if λ

is an ℓ-equilibrium for some $\tilde{\lambda}$, then $\text{st}(\lambda(a)) > 0$ only if a is sequentially optimal, hence λ is also an ℓ-equilibrium.

Characterization in terms of sequences

In this section, we characterize ℓ-equilibria for a given ℓ-tremble in terms of sequences of strategy profiles. Doing so is illustrative of the interpretation of objects defined with the use of ℓ-numbers as limits of standard objects. It will also help us clarify the relationship between ℓ-stability and other stability concepts.

Definition 4.3. Fix some ℓ-tremble $\tilde{\lambda} \in \tilde{\Lambda}$, with $\tilde{\lambda}(a) = \tilde{x}(a)\,\varepsilon^{\tilde{y}(a)}$ for all $a \in A$. The *tremble sequence associated with* $\tilde{\lambda}$, denoted $(\eta_n^{\tilde{\lambda}})_n$, is defined as $\eta_n^{\tilde{\lambda}}(a) \equiv \tilde{x}(a)\,(1/n)^{\tilde{y}(a)}$ for all $a \in A$.

We will use $G(\eta_n^{\tilde{\lambda}})$ to denote the perturbed game where each action a has to be played with probability at least $\eta_n^{\tilde{\lambda}}(a)$. Note that our notation is consistent with the definition of $\sigma_n^{\lambda}(a)$ in (3.1), now $\tilde{y}(a) > 0$ for all $a \in A$. In particular, note that $\sigma_n^{\lambda}(a) \geq \eta_n^{\tilde{\lambda}}(a)$ if and only if $\lambda(a) \geq \tilde{\lambda}(a)$.

We now fix some $\lambda \in \Lambda$ and $\tilde{\lambda} \in \tilde{\Lambda}$ with $\lambda \geq \tilde{\lambda}$. We define the *(degree of) sequential optimiality* of σ_n^{λ} with respect to $\eta_n^{\tilde{\lambda}}$, denoted $\text{SO}_n(\lambda|\tilde{\lambda})$ or $\text{SO}(\sigma_n^{\lambda}|\eta_n^{\tilde{\lambda}})$, as the maximum payoff loss generated by an action played with a probability above the trembling probability, that is,

$$\text{SO}_n(\lambda|\tilde{\lambda}) \equiv \max_{a|\sigma_n^{\lambda}(a) > \eta_n^{\tilde{\lambda}}(a)} \; \max_{a' \in A^{I^a}} \left(u(a'|\sigma_n^{\lambda}) - u(a|\sigma_n^{\lambda}) \right), \tag{4.2}$$

where $u(a'|\sigma_n^{\lambda})$ is player $\iota(I^a)$ payoff from playing a under σ_n^{λ} conditional on I^a being reached (it is uniquely defined since σ_n^{λ} has full support). Note that $\text{SO}_n(\lambda|\tilde{\lambda}) \geq 0$, and that lower values of $\text{SO}_n(\lambda|\tilde{\lambda})$ indicate that σ_n^{λ} is "more" sequentially rational given η_n^{λ}. For example, $\text{SO}_n(\lambda|\tilde{\lambda}) = 0$ if and only if σ_n^{λ} is a Nash equilibrium of $G(\eta_n^{\tilde{\lambda}})$. Similarly, defining $\varepsilon_n \equiv \text{SO}_n(\lambda|\tilde{\lambda})$, we have that σ_n^{λ} is a perfect ε_n-equilibrium of $G(\eta_n^{\tilde{\lambda}})$ (see Dilmé, 2022b). We can now provide a characterization of ℓ-equilibrium for a given ℓ-tremble in terms of the corresponding sequences.

Proposition 4.2. *λ is an ℓ-equilibrium for $\tilde{\lambda}$ if and only if $\lambda \geq \tilde{\lambda}$ and $\text{SO}_n(\lambda|\tilde{\lambda}) \to 0$ as $n \to \infty$.*

Proposition 4.2 establishes that λ is an ℓ-equilibrium for $\tilde{\lambda}$ only if σ_n^{λ} is asymptotically sequentially optimal given $\eta_n^{\tilde{\lambda}}$. Hence, assessments generated by ℓ-equilibria for $\tilde{\lambda}$ are those which are supported by sequences of strategy profiles of the perturbed games $G(\eta_n^{\tilde{\lambda}})$ which are asymptotically sequentially optimal. Intuitively, since $u(a|\sigma_n^{\lambda}) \to u(a|\lambda)$ as $n \to 0$, the requirement of sequential optimality for actions with $\lambda(a) > \tilde{\lambda}(a)$ (from Definition 4.2) is equivalent to requiring asymptotic sequential optimality for actions a with $\sigma_n^{\lambda}(a) > \eta_n^{\tilde{\lambda}}(a)$.[9]

It follows from Proposition 4.1 that λ is an ℓ-equilibrium if and only if $\mathrm{SO}_n(\lambda|\tilde{\lambda}) \to 0$ as $n \to \infty$ *for some* $\tilde{\lambda} \leq \lambda$. That is, while ℓ-equilibrium requires robustness to some ℓ-tremble, ℓ-equilibrium for a given ℓ-tremble requires robustness to that particular ℓ-tremble. As we shall see, the concept of ℓ-stable outcome will require robustness to all ℓ-trembles, hence it will be more restrictive.

4.2 ℓ-Stable outcomes

We now define the concept of ℓ-stable outcome, show its existence for any game and characterize its relationship to sequential stable outcomes. Below, we provide some properties of ℓ-stable outcomes and illustrate how they can be used to prove a given outcome is sequential stable or stable.

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

Definition 4.4. ω is *ℓ-stable* if, for each ℓ-tremble $\tilde{\lambda}$, there is an ℓ-equilibrium for $\tilde{\lambda}$ generating ω.

To briefly put our results so far into perspective, let Ω^* denote the set of ℓ-equilibrium outcomes, $\Omega^*(\tilde{\lambda})$ the set of ℓ-equilibrium outcomes for $\tilde{\lambda}$, and $\Omega^{\ell *}$ the set of ℓ-stable outcomes. Then, we have

$$\Omega^* = \bigcup_{\tilde{\lambda} \in \tilde{\Lambda}} \Omega^*(\tilde{\lambda}) \supseteq \bigcap_{\tilde{\lambda} \in \tilde{\Lambda}} \Omega^*(\tilde{\lambda}) = \Omega^{\ell *} .$$

That is, while the set of ℓ-equilibrium outcomes is the union of the sets of outcomes of ℓ-equilibria for each of the ℓ-trembles. the set of ℓ-stable outcomes is instead the intersection of the sets of outcomes of ℓ-equilibria for each of the ℓ-trembles. Since Ω^* is the set of outcomes of sequential equilibria (by Proposition 3.2), an ℓ-stable outcome is the outcome of a sequential equilibrium.

Relationsip between ℓ-stability and sequential stability

The concept of ℓ-stable outcome is similar to that of sequential stable outcome (Dilmé, 2022b). It selects behavior that remains equilibrium behavior no matter which is the form of small mistakes made

[9] Recall that perfect equilibrium requires exact optimality instead of asymptotic otpimality. Hence, if $\mathrm{SO}_n(\lambda|\tilde{\lambda}) = 0$ for all n, then σ^{λ} is a perfect equilibrium.

by the players. As we shall see, it is much simpler to use. The reason is that our definition is stated directly at the limit: It requires that each ℓ-tremble has a corresponding ℓ-equilibrium with outcome ω. Sequential stability of ω requires that all tremble sequences have corresponding sequences of perfect ε_n-equilibria (for some sequence $\varepsilon_n \to 0$) with convergent sequences of corresponding equilibrium outcomes converging to ω.

Recall that Dilmé (2022b) says that ω is a *sequential stable outcome* if, for all tremble sequences $(\eta_n)_n$, there is a sequence $(\sigma_n \geq \eta_n)_n$ such that $\omega^{\sigma_n} \to \omega$ and $\text{SO}(\sigma_n, \eta_n) \to 0$. He also defines ω to be an *extensive-form stable outcome* by requiring, instead, $\text{SO}(\sigma_n, \eta_n) = 0$ for all n. Since extensive-form stable outcomes are also outcomes of the agent-normal form of the game, they exist for games with generic payoffs. Dilmé shows that all games have a sequential stable outcome and that when an extensive-form stable outcome exists, it is sequential stable. The following result establishes that sequential stable outcomes are ℓ-stable, and hence ℓ-stable outcomes exist in all games, even in those with non-generic payoffs.

Proposition 4.3. *If ω is sequential stable, then it is ℓ-stable. Hence, an ℓ-stable outcome exists.*

The result is intuitive: Using Proposition 4.2, ℓ-stability can be thought as requiring stability with respect to the set of tremble sequences in the set $\{(\eta_n^{\tilde{\lambda}})_n | \tilde{\lambda} \in \tilde{\Lambda}\}$, while sequential stability requires stability with respect to all tremble sequences. The result is nevertheless not trivial, since sequential stability ensures the existence of some asymptotically optimal sequence $(\sigma_n)_n$ for each tremble sequence in $\{\eta_n^{\tilde{\lambda}} | \tilde{\lambda} \in \tilde{\Lambda}\}$, but this sequence need not be of the form $(\sigma_n^{\lambda})_n$ for some $\lambda \geq \tilde{\lambda}$. The proof uses the procedure in the proof of Proposition 3.1 to obain some $\lambda \in \Lambda$ satisfying that σ_n^{λ} supports $(\sigma^{\lambda}, \mu^{\lambda})$ and satisfies that $\lambda(a) \geq \tilde{\lambda}(a)$ whenever a is sequentially suboptimal. As a result, because sequential stable outcomes always exist, ℓ-stable outcomes always exist as well. Also, since stable outcomes of the agent normal form of G are sequential stable, they are also ℓ-stable.

Whether all ℓ-stable outcomes are sequential stable is an open question. The difficulty on proving this. Such a converse result is nevertheless not needed when there is a unique ℓ-stable outcome. If there is a unique ℓ-stable outcome, then it is the unique sequential stable outcome, and hence (by Dilmé, 2022b) it is the unique stable outcome of the agent normal form of G.

4.3 Obtaining ℓ-stable outcomes

The goal of this section is to provide some procedures and tools that simplify obtaining ℓ-stable outcomes. We briefly explain some strategies an economist can use to obtain ℓ-stable outcomes (these are similar to those provided in Section 4 of Dilmé, 2022b).

Necessary conditions: A first useful way to prove that an outcome ω is ℓ-stable is by ruling out the ℓ-stability alternative candidates. Then, if one rules out the ℓ-stability of all outcomes except for one, Proposition 4.3 implies that the remaining candidate is ℓ-stable.

Notably, ℓ-stable outcomes share most properties of sequential stable outcomes (see Dilmé, 2022b). For example, ℓ-stable outcomes are sequentially rational, that is, are the outcome of a sequential equilibrium. Hence, an outcome which is not the outcome of a sequential equilibrium it is not ℓ-stable, and a game with a unique sequential equilibrium outcome has a unique ℓ-stable outcome. We now highlight (and formally prove) that ℓ-stable outcomes satisfy the *never a weak best response* (NWBR) condition, which implies forward induction and iterated strict dominance (as defined in Dilmé, 2022b).

Proposition 4.4 (Never a weak best response, NWBR)**.** *Let ω be an ℓ-stable outcome. If $a \in A$ is not sequentially optimal under any ℓ-equilibrium with outcome ω, then ω is an ℓ-stable outcome of the game where a is eliminated (and all histories following it).*

It also follows from an adaptation of Dilmé's arguments that, like sequential stable outcomes, ℓ-stable outcomes satisfy on-path subgame perfection and invariance to reordering simultaneous moves.

Remark 4.1 (Signaling games). Dilmé (2022b) studies sequential stability in signaling games. He shows that sequential stable outcomes pass the classical selection (Intuitive Criterion, D1, D2) and that if there is a unique sequential stable outcome, it is stable as well. These properties also hold for ℓ-stable outcomes. The second property follows from Proposition 4.3: if a signaling game has a unique ℓ-stable outcome, it is the unique sequential stable outcome, and hence it is the unique stable outcome. We omit the proof of the first property.

Example 4.1 (Continuation of Example 3.3). As Dilmé (2022b) points out, NWBR can be used to prove that the beer outcome (where player 1 chooses $\mathrm{q_s}$ and $\mathrm{q_w}$ and player 2 chooses $\mathrm{n_q}$) is *not* ℓ-stable. Indeed, action $\mathrm{b_w}$ is never optimal under a sequential equilibrium with such outcome, so can be eliminated. In the resulting game, $\mathrm{f_b}$ is never optimal, hence can be eliminated as well. In the new game, playing $\mathrm{b_s}$ is a profitable deviation by player 1, hence there is no equilibrium with the beer outcome.

Through an ℓ-tremble: Another procedure consists in proving that a given outcome is the unique ℓ-stable outcome (or severely reducing the set of candidates to ℓ-stable outcome) by analyzing the outcomes of ℓ-equilibria for a given ℓ-tremble. If there is an ℓ-tremble $\tilde{\lambda}$ such that there is a unique ℓ-equilibrium outcome for $\tilde{\lambda}$, then such an outcome is the unique ℓ-stable outcome.

Examples 4.3-4.4 illustrate the procedure. In each case, we provide an ℓ-tremble for which there is a unique ℓ-stable outcome. In signaling games satisfying single-crossing, such ℓ-tremble typically involves "high types" (the types other types want to imitate) trembling with a higher likelihood than low types.[10] Games not satisfying a single-crossing condition may have multiple ℓ-stable outcomes. Nevertheless, through one or more ℓ-trembles, the economist may be able to reduce the set of outcome candidates to some set $\Omega^\dagger$. If all outcomes in $\Omega^\dagger$ 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.

Direct proof: The last procedure consists in proving that a given outcome is stable against any ℓ-tremble. Even though such a procedure is, in general, more difficult to apply, it is nevertheless vastly simpler to do using ℓ-numbers, where it has to be shown that for all ℓ-tremble there is an ℓ-equilibrium with this outcome, instead of showing that for any sequence of trembles there is a corresponding sequence of equilibria with outcomes converging to the desired outcome.

Example 4.2 (Continuation of Example 4.1). We now prove that the beer outcome (where player 1 chooses $\mathrm{b_w}$ and $\mathrm{b_s}$ and player 2 chooses $\mathrm{n_b}$) is ℓ-stable by explicitly showing it is stable against any ℓ-tremble (note that, in this simple example, the fact that the beer outcome is the only outcome of a sequential equilibrium other than the quiche outcome implies that the beer outcome is ℓ-stable). To see this, fix some $\tilde{\lambda} \in \tilde{\Lambda}$. It is then easy to verify that the following ℓ-strategy profile λ is an ℓ-equilibrium for $\tilde{\lambda}$ with the beer outcome. We set $\lambda(\mathrm{w}) = 1 - \lambda(\mathrm{s}) = 0.1$ and $\lambda(a) = \tilde{\lambda}(a)$ for $a \in \{\mathrm{q_s}, \mathrm{f_b}\}$. If $\tilde{\lambda}(\mathrm{q_w}) > 9\,\lambda(\mathrm{q_s})$ then $\lambda(\mathrm{q_w}) = \tilde{\lambda}(\mathrm{q_w})$ and $\lambda(\mathrm{n_q}) = \tilde{\lambda}(\mathrm{n_q})$. Otherwise, $\lambda(\mathrm{q_w}) = 9\,\tilde{\lambda}(\mathrm{q_s})$ and $\lambda(\mathrm{n_q}) = 1/2$. The rest of the actions are determined by the requirement that $\sum_{a \in A^I} \lambda(a) = 1$.

4.4 Extended ℓ-strategy profiles

In this section, we first extend the space of ℓ-numbers by including the number 0. We then extend the definitions of ℓ-strategy profile, ℓ-tremble, and ℓ-stable outcome to allow some actions to be played with probability 0. We show that all previous results remain the same when the extended concepts are used. We finally argue using extended concepts simplifies argument in practice.

Extended ℓ-numbers: We use L_0 to denote $L \cup \{0\}$, that is, the set of ℓ-numbers plus number 0. It is useful to think 0 as an ℓ-number $x\,\varepsilon^y$ with $x > 0$ and $y = \infty$, so all ℓ-numbers of the form $x\,\varepsilon^\infty$

[10] Static signaling games á la Spence (1973) satisfying a "single-crossing" condition typically have a unique stable outcome, which is the least-costly, fully-separating outcome, called the *Riley outcome* (Riley, 1979).

represent the same element (that is, for example, $2\,\varepsilon^{\infty} = 7.1\,\varepsilon^{\infty} = 0$) and the operations defined in Definition 2.1 apply (except that we do not allow dividing 0 by 0).

Extended ℓ-strategy profiles: We now extend the definition of ℓ-strategy profile (Definition 3.1) to allow actions to be played with probability 0, but requiring that information sets still receive non-zero likelihood.

Definition 4.5. An *extended ℓ-strategy profile* is a map $\lambda : A \to L_0$ satisfying that, for each $I \in \mathcal{I}$, (i) $\sum_{a\in A^I} \lambda(a) = 1$, (ii) $\lambda(a) = \pi(a)$ for all $a \in A^I$ whenever $\iota(I) = 0$, and (iii) there is some history $(a_j)_{j=1}^J \in I$ such that $\lambda(a_j) > 0$ for all $j = 1, ..., J$.

We let Λ_0 be the set of extended ℓ-strategy profiles, and $\tilde{\Lambda}_0$ be the set of extended ℓ-trembles (defined analogously). Note that while we now allow some actions to be played with zero likelihood, we require that all information sets have a history with non-zero likelihood, so conditional likelihoods can still be computed.[11] Hence, extended ℓ-strategy profiles make it simple to compute the beliefs over histories $(a_j)_{j=1}^J$ with $\lambda(a_j) = 0$ for some a_j, since then $\mu^{\lambda}((a_j)_{j=1}^J) = 0$. Extended ℓ-equilibria and extended ℓ-stable outcomes are also defined analogously. Note also that, for a given extended ℓ-equilibrium λ, we can construct an equivalent extended ℓ-equilibrium by replacing $\lambda(a)$ by $\text{st}(\lambda(a))$ for all actions a that end the game, while keeping the likelihood of all other actions the same.

Proposition 4.5. *All results in Sections 3 and 4 hold if Λ and $\tilde{\Lambda}$ are replaced by Λ_0 and $\tilde{\Lambda}_0$, respectively.*

Proposition 4.5 is useful as it permits choosing whether to work with ℓ-numbers or extended ℓ-numbers to characterize sequential equilibria and ℓ-stable outcomes. In particular, the set of ℓ-stable outcomes is the same independently of whether $\tilde{\Lambda}$ or $\tilde{\Lambda}_0$ (and whether $\Lambda^*(\tilde{\Lambda})$ or $\Lambda_0^*(\tilde{\Lambda}_0)$) is used in Definition 4.4. Working with extended ℓ-strategy profiles is particularly useful when showing that a given assessment is consistent, as it requires finding *some* ℓ-equilibrium that generates it, so having additional, easy-to-use ℓ-strategy profiles is preferable. Similarly, when ruling out the ℓ-stability of some outcome ω, it is enough to find *some* extended ℓ-tremble for which there is no ℓ-equilibrium with outcome ω. On the contrary, it is easier to prove a given assessment is *not* consistent by showing it is not generated by any ℓ-strategy profile than by showing that it is not generated by any extended ℓ-strategy profile. Similarly, proving that an outcome is ℓ-stable is easier when considering the set of ℓ-trembles instead of the set of extended ℓ-trembles.

[11] When using extended ℓ-strategy profile, equation (3.2) is replaced by $u(a|\lambda) \equiv \sum_{(a_j)_{j=1}^J \in T^a} \text{st}\Big(\frac{\prod_{a_j\neq a} \lambda(a_j)}{\lambda(I^a)}\Big)\, u_{\iota(I^a)}((a_j)_{j=1}^J)$ to accommodate for the possibility that $\lambda(a) = 0$.

	h	ℓ	$\emptyset$
H	$6,4$	$3,3$	$0,0$
L	$5,-5$	$3,-2$	$0,0$

Table 1

Example 4.3 (Continuation of Example 4.2). We now argue again that the quiche outcome is *not* ℓ-stable. Consider the extended ℓ-tremble $\tilde{\lambda}(\mathrm{b_s}) = \tilde{\lambda}(\mathrm{q_s}) = \varepsilon$ and $\tilde{\lambda}(a) = 0$ otherwise, and assume there is an extended ℓ-equilibrium λ for $\tilde{\lambda}$. If $\lambda(\mathrm{f_b}) < 0.5$ then player 1 has a profitable deviation by choosing $\mathrm{b_s}$, so it must be $\lambda(\mathrm{f_b}) \geq 0.5$. Nevertheless, for $\mathrm{f_b}$ to be optimal for player 2, it must be that $\lambda(\mathrm{b_w}) \geq 9\,\lambda(\mathrm{b_s}) > 0 = \tilde{\lambda}(\mathrm{b_w})$, so $\mathrm{b_w}$ must be optimal for player 1. This is a contradiction because player 1 obtains 3 by playing $\mathrm{q_w}$ and at most 2 by playing $\mathrm{b_w}$.

4.5 Example: dynamic signaling

Example 4.4 (Dynamic signaling). The refinements for static signaling games (e.g., Intuitive Criterion, D1, D2) cannot be directly applied to dynamic signaling games. As a result, ad hoc selection criteria are often used. In this example, we illustrate how ℓ-stability has the potential to select equilibria across different models using the same criterion, hence easing the comparison of the results. We do so by analyzing, in a simple and unified framework, dynamic signaling models with public (Noldeke and Van Damme, 1990) and private (Swinkels, 1999) with preemptive offers (see also Hörner and Vieille, 2009). This example illustrates the simplicity of using ℓ-numbers to make stability arguments.

Consider the following dynamic signaling problem. There are two periods. There is a different firm in every period (firm 1 and firm 2). There is one student who has one of two types, $\theta = L, H$. The prior about $\theta = H$ is $1/2$. Only the student knows her type. Each firm has 3 options: make a high offer (h), make a low offer (ℓ), or pass ($\emptyset$). If an offer is made, the student either accepts it, A, or rejects it, R. If acceptance happens in the first period, the game ends; otherwise, the student pays a signaling cost and moves to the second period. The signaling cost c_θ is type-dependent, and satisfies $0 < c_H < c_L < 1$. The payoffs (without education costs) are given in Table 1. Notice that the expected payoff for firm 1 if all types of the student accept for sure a low offer is $1/2$, while if it offers h, its expected payoff is $-1/2$. Therefore, firm 1 never offers h since it would be accepted for sure by the students and gives it a negative payoff.

1. *Public offers case:* In the public offers case, the second firm observes the first firm's offer. There

are two outcomes of sequential equilibria: one where firm 1 offers ℓ and it is accepted for sure by the student, and one where firm 1 offers $\emptyset$ and firm 2 offers ℓ which is accepted for sure. We now prove that there is a unique ℓ-stable outcome.

Fix an ℓ-stable outcome ω, some $\tilde{\lambda}$ and some $\lambda \in \Lambda^*(\tilde{\lambda})$ with $\omega^\lambda = \omega$. We use the notation that $R_{\theta a}$ is the action of rejection by the θ after $a \in \{\ell, h\}$ has been offered in the first period. Assume $\tilde{y}(R_{L\ell}) > \tilde{y}(R_{H\ell})$, that is, the H-student trembles to rejecting ℓ from firm 1 more often than the L-student. Assume that firm 1 offers ℓ with positive probability. Such an offer must be accepted with positive probability by the H-student, which implies that the L-student has a strict incentive to accept it, and hence

$$y(R_{L\ell}) = \tilde{y}(R_{L\ell}) > \tilde{y}(R_{H\ell}) \geq y(R_{H\ell}) \ .$$

Still, if $y(R_{L\ell}) > y(R_{H\ell})$, then necessarily firm 2 offers h, which leads to a contradiction. So, in the unique ℓ-stable outcome, firm 1 offers $\emptyset$ and firm 2 offers ℓ, which is accepted by both student types. Therefore, ℓ-stability (and sequential stability) does not always select the least costly, maximally separating equilibria.

2. *Private offers case:* In the private offers case, firm 2 does not observe the firm 1's offer. There are again two outcomes of sequential equilibria: one where firm 1 offers ℓ and it is accepted for sure by the student, and another outcome is described below. We now prove that there is a unique ℓ-stable outcome.

   Fix an ℓ-stable outcome ω, some $\tilde{\lambda}$ and some $\lambda \in \Lambda^*(\tilde{\lambda})$ with $\omega^\lambda = \omega$. Assume that $\tilde{y}(R_{H\ell})$ is the unique minimum of $\{\tilde{y}(a) | a \in A\}$. Firm 1 makes a low offer and the student accepts it with positive probability, since otherwise firm 2 would offer ℓ in equilibrium, but then firm 1 would profitably deviate to offer ℓ as it would be accepted for sure. Also, firm 1 cannot hire with probability one since, otherwise, a rejection in the period would lead firm 2 to believe the type is H for sure and would offer h, so the student would have an incentive to deviate.[12] In the only option left, firm 1 randomizes between ℓ (which is accepted for sure by L and with interior probability by H) and $\emptyset$, and firm 2 randomizes between ℓ and h, both accepted for sure.

[12]Formally, assume that firm 1 hires with probability one. An off-path rejection in the first period could come from a deviation by firm 1 (by offering $\emptyset$) or from a deviation of the student (to reject). Since firm 1 earns less by offering $\emptyset$ than by offering ℓ, we have that firm 2 deems a deviation by the student as infinitely more likely than one by firm 1. Also, if $y(R_{L\ell}) \leq \tilde{y}(R_{H\ell}) < y(R_{L\ell})$ it must be that L is indifferent on rejecting ℓ, but then H is strictly willing to reject. It must then be that $y(R_{L\ell}) > \tilde{y}(R_{H\ell}) \geq y(R_{H\ell})$, but this implies that firm 2 offers h for sure, hence we have a contradiction.

5 Conclusions

In our view, this paper has three contributions. The first is developing a new language to analyze tremble-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.

The final contribution is to provide new equilibrium concepts using the language of ℓ-numbers. Most saliently, we define ℓ-stable outcomes, 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 sequential stable when it is the unique ℓ-stable outcome, our analysis provides a method to obtain sequential 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*, 55, 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.

Dilmé, F., 2022a, "A characterization of consistent assessments using power sequences of strategy profiles," mimeo.

——— , 2022b, "Sequential stable outcomes," mimeo.

Fudenberg, D., and J. Tirole, 1991, "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.

Henle, J. M., and E. M. Kleinberg, 2014, *Infinitesimal calculus*. Courier Corporation.

Hörner, J., and N. Vieille, 2009, "Public vs. private offers in the market for lemons," *Econometrica: Journal of the Econometric Society*, 77(1), 29–69.

Keisler, H. J., 2013, *Elementary calculus: An infinitesimal approach*. Courier Corporation.

Kohlberg, E., and J.-F. Mertens, 1986, "On the strategic stability of equilibria," *Econometrica: Journal of the Econometric Society*, 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*, 50(4), 863–894.

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., 1978, "Refinements of the Nash equilibrium concept," *International Journal of Game Theory*, 7(2), 73–80.

——— , 1986, "Multistage games with communication," *Econometrica: Journal of the Econometric Society*, pp. 323–358.

Noldeke, G., and E. Van Damme, 1990, "Signalling in a dynamic labour market," *The Review of Economic Studies*, 57(1), 1–23.

Riley, J. G., 1979, "Informational equilibrium," *Econometrica: Journal of the Econometric Society*, 47, 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.

Swinkels, J. M., 1999, "Education signalling with preemptive offers," *The Review of Economic Studies*, 66(4), 949–970.

A Omitted Proofs

A.1 Proofs of the Results in Section 3

Proof of the Proposition 3.1

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

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 $y(a) = 0$, 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 $y(a) = 0$ 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. We divide the proof into four steps:

Step 2. "Only if" part of the statement of Proposition 4.2. We begin by proving that, if λ is an ℓ-equilibrium for $\tilde{\lambda}$, then $\mathrm{SO}_n(\lambda|\tilde{\lambda}) \to 0$ as $n \to \infty$. Since $\lambda(a) \geq \tilde{\lambda}(a)$ we have that $\sigma_n^\lambda(a) \geq \eta_n^{\tilde{\lambda}}(a)$ for all n. Furthermore, $\sigma_n^\lambda(a) > \eta_n^{\tilde{\lambda}}(a)$ only if $\lambda(a) \geq \tilde{\lambda}(a)$, that is, only if

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

since λ is an ℓ-equilibrium for $\tilde{\lambda}$ (see Definition 4.2). Since $u(a|\lambda) = \lim_{n\to\infty} u(a|\sigma_n^\lambda)$, we have that

$$\lim_{n\to\infty} \max_{a \in A^{I^a}} \{0, u(a'|\sigma_n^\lambda) - u(a|\sigma_n^\lambda)\} = 0 \, .$$

By letting A^* be the set of actions a with $\lambda(a)>\tilde{\lambda}(a)$, we define

$$\varepsilon_n \equiv \max_{a\in A^*}\max_{a\in A^{I^a}} \left\{0, u(a'|\sigma_n^\lambda)-u(a|\sigma_n^\lambda)\right\} .$$

Note that $\varepsilon_n \to 0$. Furthermore, we have that $u(a|\sigma_n^\lambda)\geq u(a'|\sigma_n^\lambda)-\varepsilon_n$ for all $a\in A^*$, that is, $\mathrm{SO}_n(\lambda|\tilde{\lambda})\in [0,\varepsilon_n]$. This proves the desired result.

Step 2. "If" part of the statement of Proposition 4.2. We now prove that, if λ is such that $\lambda\geq\tilde{\lambda}$ and $\mathrm{SO}_n(\lambda|\tilde{\lambda})\to 0$ as $n\to\infty$, then λ is an ℓ-equilibrium for $\tilde{\lambda}$. Take some action a such that $\lambda(a)>\tilde{\lambda}(a)$. Note that $\mathrm{SO}_n(\lambda|\tilde{\lambda})\to 0$ implies that

$$\lim_{n\to\infty}\max_{a'\in A^{I^a}}\left(u(a'|\sigma_n^\lambda)-u(a|\sigma_n^\lambda)\right)=0 .$$

This implies that $u(a'|\lambda)\leq u(a|\sigma^\lambda)$ for all $a'\in A^{I^a}$. Hence, since $\lambda(a)>\lambda(a)$ only if a is sequentially rational (according to λ), it follows that λ is an ℓ-equilibrium for $\tilde{\lambda}$. □

Proof of Proposition 4.3

Proof. Let ω be a sequential stable outcome, and let $\tilde{\lambda}$ be an ℓ-tremble. Since ω be sequential stable, there exists is a sequence $(\sigma_n\geq\eta_n^{\tilde{\lambda}}\})_n$ such that $\omega^{\sigma_n}\to\omega$ and $\mathrm{SO}(\sigma_n,\eta_n)\to 0$. We now want to prove that there exists an ℓ-equilibrium for $\tilde{\lambda}$ with outcome ω. Taking a subsequence if necessary, we assume that σ_n supports some assessment (σ,μ).

We now recall a result from the proof of Theorem 3.1 in Dilmé (2022a).

Lemma A.1 (Lemma 3.1 in Dilmé, 2022a)**.** *Let $(\sigma_n : A\to(0,1])_n$ be a sequence. 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 are unique $\alpha^0(a)\in\mathbb{R}$ and $\alpha(a)\equiv(\alpha^1(a),...,\alpha^K(a))\in\mathbb{R}^K$ such that*

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

(Proof of Proposition 4.3 continues.)

We apply Lemma A.1 to $(\sigma_n)_n$, and we let $(j_n\in\mathbb{N})_n$, $((q_n^1,...,q_n^K)\in\mathbb{R}_{++}^K)_n$, and $(\alpha_0,\alpha):A\to\mathbb{R}^{K+1}$ be the corresponding sequences and functions. We let λ be an ℓ-equilibrium such that $(\sigma^\lambda,\mu^\lambda)=(\sigma,\mu)$ (which exists by Proposition 3.2). Dilmé's analysis implies that there is some large enough $M\in\mathbb{R}_{++}$ satisfying that

$$\lambda(a)\equiv\alpha^0(a)\,\varepsilon^{K\sum_{k=1}^K M^k\,\alpha^k(a)} .$$

is such that $(\sigma^\lambda, \mu^\lambda) = (\sigma, \mu)$. We let $\hat{A}$ be the set of actions sequentially optimal under λ (and hence under (σ, μ)). There are two cases:

1. Assume first $\hat{A} = A$. In this case, all actions are sequentially optimal under (σ, μ). If $\lambda(a) \geq \tilde{\lambda}(a)$ for all $a \in A$, then λ is an ℓ-equilibrium for $\tilde{\lambda}$ with outcome ω. Otherwise, for all $a \in A$, define $\hat{\lambda}(a)$ as

$$\hat{\lambda}(a) = x(a)\, \varepsilon^{K\, y(a)}$$

for K large enough that $\hat{\lambda}(a) \geq \tilde{\lambda}(a)$ for all $a \in A$. It is easy to verify that $(\sigma^{\hat{\lambda}}, \mu^{\hat{\lambda}}) = (\sigma, \mu)$. Hence, $\tilde{\lambda}$ is a ℓ-equilibrium for $\tilde{\lambda}$ with outcome ω.

2. Assume then that $\hat{A} \neq A$, and fix some $\hat{a} \in A \backslash \hat{A}$. We first note that it must be that, for all $a' \in \hat{A} \neq A$, we have

$$(\alpha_0(\hat{a})/\tilde{x}(\hat{a}))^{1/\tilde{y}(\hat{a})} = (\alpha_0(\hat{a}')/\tilde{x}(\hat{a}'))^{1/\tilde{y}(\hat{a}')} \text{ and } \alpha(\hat{a})/\tilde{y}(\hat{a}) = \alpha(\hat{a}')/\tilde{y}(\hat{a}')\,,$$

or equivalently

$$\tilde{x}(\hat{a}') = \left(\left(\tfrac{\tilde{x}(\hat{a})}{x(\hat{a})}\right)^{1/y(\hat{a})}\right)^{y(\hat{a}')} x(a') \text{ and } \tilde{y}(\hat{a}') = \tfrac{\tilde{y}(\hat{a})}{y(\hat{a})}\, y(\hat{a}')\,. \tag{A.2}$$

This follows from part 2 of Lemma A.1, and the fact that

$$\frac{(\sigma_n(\hat{a})/\tilde{x}(\hat{a}))^{1/\tilde{y}(\hat{a})}}{(\sigma_n(\hat{a}')/\tilde{x}(\hat{a}'))^{1/\tilde{y}(\hat{a}')}} = \frac{(\eta_n^{\tilde{\lambda}}(\hat{a})/\tilde{x}(\hat{a}))^{1/\tilde{y}(\hat{a})}}{(\eta_n^{\tilde{\lambda}}(\hat{a}')/\tilde{x}(\hat{a}'))^{1/\tilde{y}(\hat{a}')}} = \frac{n^{-1}}{n^{-1}} = 1\,.$$

Define $\mu \equiv \left(\frac{\tilde{x}(\hat{a})}{x(\hat{a})}\right)^{1/y(\hat{a})}$ and $\nu \equiv \frac{\tilde{y}(\hat{a})}{y(\hat{a})}$. Define also, for all $a \in A$,

$$\hat{\lambda}(a) = \mu^{\lambda(a)}\, x(a)\, \varepsilon^{\nu\, y(a)}\,.$$

It is easy to see that this is an ℓ-strategy profile. It is also clear that $\hat{\lambda}(a) = \tilde{\lambda}(a)$ for all $a \in A \backslash \hat{A}$ (recall equation (A.2)). It is only left to show that $\hat{\lambda}(a) \geq \tilde{\lambda}(a)$. This follows from the fact that since, for all a, we have $\sigma_n(\hat{a}) \geq \eta_n^{\tilde{\lambda}}(\hat{a}')$, we also have

$$\frac{(\sigma_n(\hat{a})/\tilde{x}(\hat{a}))^{1/\tilde{y}(\hat{a})}}{(\sigma_n(a)/\tilde{x}(a))^{1/\tilde{y}(a)}} \leq 1\,,$$

which, from the previous arguments, implies that $\hat{\lambda}(a) \geq \tilde{\lambda}(a)$.

□

A.3 Proofs of the Results in Section 4.3

Proof of Proposition 4.4

Proof. Let ω be a sequential stable outcome. Let $\hat{a} \in A^I$ be an action that is *not* sequentially optimal in any sequential equilibrium with outcome ω. 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'. Define an ℓ-tremble $\tilde{\lambda}$ of G as follows:

$$\tilde{\lambda}(a) \equiv \begin{cases} \tilde{\lambda}(a) & \text{if } a \in A', \\ \prod_{a' \in A'} \tilde{\lambda}(a') & \text{otherwise,} \end{cases}$$

for all $a \in A$. Note that $\tilde{\lambda}$ is such that any history not belonging to G' (i.e., with some $a \notin A'$) has zero relative likelihood with respect to any history of G'. Let λ be an ℓ-equilibrium for $\tilde{\lambda}$ with outcome ω (which exists since ω is ℓ-stable). It is then clear that the restriction of λ to A' is an ℓ-equilibrium for $\tilde{\lambda}$. Indeed, $\hat{a}$ is not sequentially optimal under λ because it is not sequentially optimal under $(\sigma^\lambda, \mu^\lambda)$, since $(\sigma^\lambda, \mu^\lambda)$ is a sequential equilibrium.. □

Proof of Proposition 4.5

Proof. The general idea behind the result is that any result about a extended ℓ-strategy profile/ℓ-tremble can be proven reproducing the proof of the analogous result using a ℓ-strategy profile/ℓ-tremble assigning very small likelihoods to actions assigned likelihood 0.

1. Proving Proposition 3.1 for extended ℓ-strategy profiles is straightforward. It is clear that since $\Lambda_0 \supset \Lambda$, each consistent assessment is generated by some extended ℓ-strategy profile. We then prove that each extended ℓ-strategy profile λ_0 generates a consistent assessment. Define, for all $a \in A$,

$$\lambda(a) \equiv \begin{cases} \lambda_0(a) & \text{if } \lambda_0(a) \neq 0, \\ \varepsilon^K & \text{if } \lambda_0(a) = 0, \end{cases}$$

   where $K \equiv |A|\,(1 + \max\{y_0(a) | \lambda_0(a) \neq 0\})$. Note that λ is an ℓ-strategy profile. It is easy to see that λ and λ_0 generate the same assessment, hence the assessment generated by λ_0 is consistent.

2. Proving Proposition 3.2 for extended ℓ-equilibrium is also straightforward. The "only if" part follows from the fact that $\Lambda_0^* \supset \Lambda$, so all sequential equilibria are generated by some extended

ℓ-equilibrium. To prove the "if" part, let λ_0 be a extended ℓ-equilibrium. Define

$$\sigma_n^\lambda(a) \equiv \begin{cases} (1/n)^K & \text{if } y(a)=\infty, \\ x(a)(1/n)^{y(a)} & \text{if } y(a)\in(0,\infty), \\ M_n(\lambda, I^a)\,x(a) & \text{if } y(a)=0, \end{cases}$$

where $K = 1+\sum_{a|y(a)\neq\infty} y(a)$. where $M_n(\lambda, I^a)$ is a factor that ensures that $\sum_{a'\in A^{I^a}} \sigma_n^\lambda(a')=1$ and where the subindex n is initialized so that $M_n(\lambda, I^a)\geq 0$ for all n and a. It is then easy to see that $(\sigma_n^\lambda)_n$ generates (μ,σ).

3. Propositions 4.1 and 4.2 (allowing that $\tilde{\lambda}\in\tilde{\Lambda}_0$ and $\lambda\in\Lambda_0$ in their statements) are proven analogously.

4. Propositions 4.3 and 4.4 follow from the fact, to be proven now, that the set of ℓ-stable outcomes coincides with the set of extended ℓ-stable outcomes (allowing $\tilde{\lambda}\in\tilde{\Lambda}_0$ and $\lambda\in\Lambda_0$ in Definition 4.4). We divide the proof into two parts:

   (a) Let ω_0 be a extended ℓ-stable outcome. Take an ℓ-tremble $\tilde{\lambda}\in\tilde{\Lambda}$. Because $\tilde{\Lambda}\subset\tilde{\Lambda}_0$ and ω_0 is extended ℓ-stable, there is a extended ℓ-equilibrium for $\tilde{\lambda}$, $\lambda_0\in\Lambda_0^*(\tilde{\lambda})$, with outcome ω_0. Since $\lambda_0(a)\geq\tilde{\lambda}(a)>0$, we have that $\lambda_0\in\Lambda^*(\tilde{\lambda})$. This proves that ω_0 is ℓ-stable.

   (b) Let ω be an ℓ-stable outcome. Take a extended ℓ-tremble $\tilde{\lambda}_0\in\tilde{\Lambda}_0$. Define, for all $a\in A$,

$$\tilde{\lambda}(a) \equiv \begin{cases} \tilde{\lambda}_0(a) & \text{if } \tilde{\lambda}_0(a)\neq 0, \\ \varepsilon^K & \text{if } \tilde{\lambda}_0(a)=0, \end{cases}$$

   where $K\equiv|A|\,(1+\max\{y_0(a)|\tilde{\lambda}_0(a)\neq 0\})$. Note that $\tilde{\lambda}\in\tilde{\Lambda}$. Because ω is ℓ-stable, there is an ℓ-equilibrium for $\tilde{\lambda}$, $\lambda\in\Lambda^*(\tilde{\lambda})$, with outcome ω. Define, for all $a\in A$,

$$\lambda_0(a) \equiv \begin{cases} \lambda(a) & \text{if } \lambda(a)>\varepsilon^K, \\ 0 & \text{if } \lambda(a)=\varepsilon^K. \end{cases}$$

   Note that $\lambda_0(a)\geq\tilde{\lambda}_0(a)$ for all $a\in A$, and that λ_0 generates the same outcome as λ. Also, note that $\lambda_0(a)>\tilde{\lambda}_0(a)$ only if $\lambda(a)>\tilde{\lambda}(a)$, which implies that $\lambda_0\in\Lambda_0^*(\tilde{\lambda}_0)$. This proves that ω is extended ℓ-stable.

□

B Relationship to CPSs and LPSs

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

B.1 Conditional Probability Systems

We first relate ℓ-strategy profiles to conditional probability systems, studied in Battigalli (1996), and to relative probability spaces, studied in Myerson (1986) and Kohlberg and 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 behavior strategies. Battigalli (1996) defines a *conditional probability system (CPS)* on S as 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 $S_1, S_2, S_3 \subset S$,

$$S_1 \subset S_2 \subset S_3 \text{ implies } P(S_1|S_3) = P(S_1|S_2)\,P(S_2|S_3). \tag{B.1}$$

To ensure that a CPS is consistent with independent randomizations, Battigalli defines the following property: a CPS P *has 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[13]

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

Finally, Battigalli defines a *strategic extended assessment* as a triple (P, σ, μ), where (σ, μ) is an assessment and P is a CPS satisfying

$$\mu^I(h) = P\big(S(h) \big| S(I)\big) \;\text{ and }\; \sigma(a) = P\Big(\{s \in S(I^a) \,|\, s(I^a) = a\} \,\Big|\, S(I^a)\Big) \tag{B.2}$$

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 $\sum_{s \in S_1} \lambda(s)$.

[13] Our definition is slightly different than Battigalli's in that we define the independence property over pure behavior 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, doing so is without loss of generality since each player randomizes independently in each of her information sets.

Finally, the *conditional probability system generated by* λ, denoted P^{λ}, as

$$P^{\lambda}(S_1|S_2) \equiv \text{st}\left(\frac{\lambda(S_1 \cap S_2)}{\lambda(S_2)}\right) \quad \text{for all } S_1, S_2 \subset S \text{ with } S_2 \neq \emptyset\ . \tag{B.3}$$

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

Proposition B.1. *For each ℓ-strategy profile λ, P^{λ} has the independence property and $(P^{\lambda}, \mu^{\lambda}, \sigma^{\lambda})$ is a strategic extended assessment.*

Corollary 3.1 in Battigalli (1996) shows that an assessment (σ, μ) is consistent only if (P, σ, μ) is a strategic extended assessment for some P with the independence property. Proposition B.1 establishes that one such P can be obtained from an ℓ-strategy profile λ generating (σ, μ). Battigalli further shows that, while the indepencence property is necessary, it is not sufficent.[14] Kohlberg and Reny (1997) obtain a similar result for relative probability systems: they show that, while a necessary condition for consistency can be written in terms of weak indepencence and coordinate-wise exchangeablilty of some random variables in a relative probability space, such condition is not sufficient. In contrast, Proposition 3.1 establishes a necessary and sufficient condition for consistency of an assessment: being generated by an ℓ-strategy profile.

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-strategy profiles, which is $2(|A|-|\mathcal{I}|)$). Furthermore, the requirements that a given CPS has 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.

B.2 Lexicographic probability systems

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

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 *order* $\hat{K} \in \mathbb{Z}_+$, such that, for all $a \in A$, there is some k such that $\sigma^k(a) > 0$. We use $\hat{K}(a)$ to denote $\min\{k|\sigma^k(a) > 0\}$.

[14]He uses the tree in Figure 3 and argues that, while a strategic extended assessment generated by P satisfying that $P(t_i|\{t_i, t_{i'}\}) = 0$ for all $i > i'$ satisfies the indepencence property, it does not satisfy consistency (as we proved in Example 3.4).

As Govindan and Klumpp point out, $(\sigma^k|_{\iota^{-1}(i)})_{k=0}^{\hat{K}}$ can be interpreted as a collection of theories of players $N\backslash i$ about player 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 $(a_j)_{j=1}^J \in H$ in some information set $I \in \mathcal{I}$, the following probability and belief:

$$\sigma(a) = \sigma^0(a) \text{ and } \mu((a)_{j=1}^J) = \begin{cases} 0 & \text{if } \hat{K}((a_j)_{j=1}^J) > \hat{K}(I), \\ C_I \prod_{j=1}^J \sigma^{\hat{K}(a_j)}(a_j) & \text{if } \hat{K}((a_j)_{j=1}^J) = \hat{K}(I), \end{cases} \tag{B.4}$$

where $\hat{K}((a_j)_{j=1}^J) = \sum_{j=1}^J \hat{K}(a_j)$, $\hat{K}(I) = \min_{h\in I} \hat{K}(h)$, and C_I is the constant that keeps $\mu|_I$ a probability distribution. In words, the strategy profile in the assessment generated by an 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.2. *Let $\hat{\sigma} \equiv (\sigma^k)_{k=0}^{\hat{K}}$ be a behavior LPS. Then, $a \mapsto (\hat{K}(a), \sigma^{\hat{K}(a)}(a))$ for all $a \in A$ is an ℓ-strategy profile that generates the same assessment as $\hat{\sigma}$.*

Proposition B.2 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. We further show a necessity result: for an assessment to be consistent, it has to be generated by some ℓ-strategy profile.

In applications, especially in extensive-form games, using LPSs to study behavior may require high values for $\hat{K}$, diminishing their usefulness. Hence, the reduced dimensionality of ℓ-strategy profiles ($\mathbb{R}^{2(|A|-|\mathcal{I}|)}$ instead of $\mathbb{R}^{(\hat{K}+1)(|A|-|\mathcal{I}|)}$ in a behavior LPS, where $\hat{K}$ is potentially a large number) and their simple additive and multiplicative properties make them easier to work in practice.[15]

Example B.1. Figure 4 provides a simple game and a (consistent) assessment that cannot be represented by an LPS with order strictly lower than 5. For such assessment, one needs to specify 30 distributions over actions (5 for each information set). Nevertheless, most of such distributions are irrelevant to assess sequential optimality.

[15] 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 **??**, the equilibrium paths of perfect and sequential equilibria coincidence for generic payoffs.

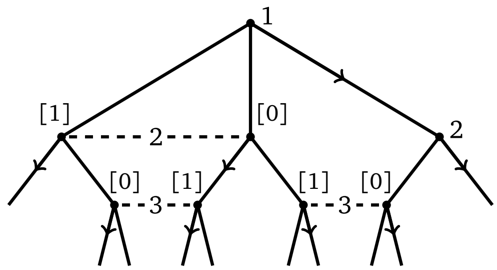

Figure 4 Figure from Example B.1.

B.3 Proofs of the Results in Section B

Proof of Proposition B.1

Proof. It is clear that P^{λ} satisfies (B.1). 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 $A^{\mathcal{I}''}$.[16] 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 $h \equiv (a_j)_{j=1}^{J}$,

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

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.3)), the first condition in equation (B.2) holds. To prove that the second condition in equation (B.2) 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

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

[16]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 $S' \subset A^{\mathcal{I}'}$.

Then, the second condition in equation (B.2) holds, and the proof is done. □

Proof of Proposition B.2

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 $\sum_{a\in A^I} \lambda(a) = \left(\sum_{a\in A^I} \sigma^0(a)\right) \varepsilon^0 = 1$.

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

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

and also that

$$\lambda(I) = \sum_{h\in I} \lambda(h) = \left(\sum_{h\in I|y(h)=y(I)} x(h)\right) \varepsilon^{y(I)} .$$

Since $y(h)=\hat{K}(h)$ and $y(I)=\hat{K}(I)$ 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 $y(h)=y(I)$), then

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

hence the result holds. □