Likelihood Assignments in Extensive-Form Games

PRELIMINARY AND INCOMPLETE

Francesc Dilmé[*]

April 3, 2017

Abstract

The paper introduces *likelihood assignments* and *perturbed likelihood assignments* in extensive-form games. They assign a *likelihood order* and a *likelihood intensity* to each action in the game, which can be interpreted as its limit likelihood in a sequence of strategy profiles or trembles. An assessment is shown to be a sequential equilibrium if and only if it is generated by a likelihood assignment which *optimally supersedes* some perturbed likelihood assignment, that is, if it assigns a higher likelihood to each action, and a strictly higher likelihood only to actions that are optimal. Similarly, an outcome is stable only if any perturbed likelihood assignment is optimally superseded by a likelihood assignment which leads to the outcome. We illustrate the use of likelihood assignments through some examples.

Key words: Likelihood assignment, sequential equilibrium, stable outcome.

JEL classification codes: C72, C73.

[*]University of Bonn. fdilme@uni-bonn.de

1 Introduction

Consistency, sequential rationality and strategic stability are plausible requisites to the behavior expected from strategic rational agents in dynamic settings. Still, in many applications, proving or disproving that equilibria or equilibrium outcomes satisfy these requirements may be far from easy. Indeed, when applied in particular dynamic games, using them implies proving or disproving the existence of sequences of strategy profiles (or trembles) satisfying some properties, which sometimes becomes an intractable problem. As a result, (less theoretically appealing but easier to use) refinements and selection criteria have been developed, many of them usable in a small subset of applications.[1]

The goal of this paper is to provide a new set of tools intended to simplify the characterization of sequential equilibria and stable outcomes. We compare them to the previously proposed tools, and illustrate their use in a number of examples.

Likelihood assignments assign, to each action in an extensive-form game, a non-negative real number indicating its likelihood order, and a strictly positive real number interpreted as its likelihood intensity. They satisfy three restrictions: (1) in each information set there are actions with likelihood order equal to 0 (interpreted as them being played with positive probability conditional on the information set being reached); (2) the sum of likelihood intensities of actions of order 0 is 1 (in this case, the likelihood intensities are interpreted as probabilities); and (3) they are consistent with nature's probabilistic choices. The likelihood assignment of a history is then obtained by adding the orders and multiplying the intensities of the actions that lead to it.

It is important to stress four features of likelihood assignments that differ from previous similar constructions (like lexicographic probability systems, see literature review below). First, a likelihood assignment is defined on actions, not pure strategies, and it imposes restrictions information set by information set. This simplifies their use in applications where the set of pure strategy profiles is large, and therefore checking whether some global consistency conditions are satisfied is, in general, difficult. Second, given the additive-multiplicative properties of likelihood assignments, computing the relative likelihood of different histories is trivial, and so is computing beliefs. Third, the likelihood order is allowed to be a non-negative real number instead of a non-negative integer. This gives flexibility to likelihood assignments in games with multiple information sets. Finally, the likelihood intensities of actions that are not of order 0 are not bounded. This feature, again, gives flexibility, and allows us to both define the the

[1]Prominent examples Are the use of perfect Bayesian equilibria in all its forms (see, for example, Fudenberg and Tirole (1991)), and also refinements and selection criteria in the (static) signaling literature, initiated by the seminal papers of Cho and Kreps (1987) and Banks and Sobel (1987).

intensities of actions in an independent way and still make the intensities of histories useful to compare their relative likelihood.

Each likelihood assignment generates a unique assessment which, in each information set, only assigns positive probability to its histories of the lowest likelihood order, with their relative likelihood given by the intensities that the likelihood assignment assigns to them. We obtain that likelihood assignments are closely related to consistent assessments: in fact, an assessment is consistent if and only if it is generated by some likelihood assignment. This result provides us with a simple procedure to prove the consistency of an assessment, as it only requires providing a likelihood assignment generating it.

In order to relate likelihood assignments to the stability of equilibrium outcomes, we introduce the concept of *perturbed likelihood assignments*: they assign, to each action in the game, two strictly positive real numbers, interpreted as the limit likelihood order and intensity of a sequence of trembles of the game. A perturbed likelihood assignment is interpreted as giving the "minimal likelihood" with which each action is played the limit of the a sequence of perturbed games. This motivates defining an *assignment equilibrium* for a given perturbed likelihood assignment as likelihood assignment that (1) *supersedes* the perturbed likelihood assignment, that is, it assigns likelihood to each action weakly higher likelihood and (2) only assigns a strictly higher likelihood to actions which maximize the continuation payoff. We show that an assessment is a sequential equilibrium if and only if it is an assignment equilibrium for some likelihood assignment.

We finally obtain a necessary condition for outcomes of extensive-form games to be stable in the sense of Kohlberg and Mertens (1986). To do this, we define an outcome to be *assignation-stable* if for any perturbed likelihood assignment there is an assignment equilibrium generating the outcome. Our concept can be interpreted as the analogous "limit" version of the concept of stable outcome. We show that a stable outcome is assignation-stable, so there is generic existence of assignation-stable outcomes, and when there is a unique assignation-stable outcome, it is stable. We again illustrate this result through examples and applications.

1.1 Related Literature

Since the definition of Nash equilibrium (Nash, 1951), many refinements (and selection criteria) have been developed, specially in extensive-form games. The main objective of these refinements is selecting equilibria which have some "plausible" properties or, equivalently, ruling out implausible behavior which takes place in some Nash equilibria. Of particular relevance is the definition of sequential equilibria Kreps and

Wilson (1982) (preceded by the concepts of perfect equilibria (Selten, 1975) and proper equilibria (Myerson, 1978)). Still, in many applications, the concept usually used is (one of the several versions of) “perfect Bayesian equilibrium” (PBE). The reason is, mainly, that computing PBEs is easier than computing sequential equilibria, as it avoids proving or disproving the existence of sequences of strategy profiles with some peroperties in order to establish the beliefs at off-the-path nodes of the game. Analogously, even though the concept of strategic stability (Kohlberg and Mertens, 1986) is appealing, its practical use is very limited for similar reasons. Our characterizations aim at simplifying the use of these concepts, and we show their usefulness through some examples.

Our approach shares some similarities to conditional probability systems, relative probabilities and lexicographic probability systems (LPSs), which have been used to model each player's belief about the strategies chosen by the rest of the players in perfect and proper equilibria,[2] and particularly to Battigalli (1996), Kohlberg and Reny (1997) and Govindan and Klumpp (2003), who apply conditional probability systems (CPSs), relative probabilities and LPSs directly to the strategies of the players. As we discuss in detail in Appendix B, a likelihood assignment can be understood as retaining only the information in a conditional probability systems or an LPS necessary to ensure the consistency of assessments and to evaluate their sequential rationality and, as a result, to characterize sequential equilibria (but not perfect or proper equilibria[3] It does so in a flexible way: the fact that the likelihood order is allowed to be a real number and the likelihood intensity is not bounded, which eases its practical use. Furthermore, it does not require global conditions (such as the “independence property” for LPSs or CPSs,) but only conditions that are sastisfied information set by information set conditions. Finally, we show that assigning likelihoods to histories (or over pure strategies) can be done in a simple and intuitive way: assigning a likelihood order/intensity to each action and adding/multiplying them accordingly, which eases the obtention consistent beliefs systems.

The rest of the paper is organized as follows: after this introduction, in Section 2 we introduce the notation for extensive form games used in the rest of the paper. Next, in Section 3, we define likelihood assignments and characterize consistent assessments.

---

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

[3]Recall that Kreps and Wilson (1982) show that the equilibrium paths of perfect and sequential equilibria coincidence for generic payoffs.

Section 4 characterizes sequential equilibrium and stable outcomes. Finally, Section 5 concludes. Appendix A contains the omitted proofs, while Appendix B contains an exposition of the relationship between likelihood assignments, LPSs and CPSs.

2 Extensive Form Games

We base our definitions on Osborne and Rubinstein (1994), although we do not follow them exactly.

Definition 2.1. A (finite) *extensive-form game* has the following components.

1. **Histories:** A finite set of *actions* A and finite set of sequences of actions T (or *histories*) satisfying:
   (a) The empty sequence $\emptyset$ is a member of T.
   (b) If $(a_j)_{j=1}^{J} \in T$ and $J' < J$ then $(a_j)_{j=1}^{J'} \in T$.

   A history $(a_j)_{j=1}^{J} \in T$ is *terminal* if there is no a_{J+1} such that $(a_j)_{j=1}^{J+1} \in T$. The set of terminal histories is denoted Z. For any $t \in T$, we use $A^t \equiv \{a|(t,a) \in T\}$ to denote the set of actions available at t (notice that $A^z = \emptyset$ for all $z \in Z$.)
2. **Information:** A partition H of the non-terminal histories, such that for each *information set* $h \in H$, if $t \in h$ and $a \in A^t$ then, for all $t' \in T$, $a \in A^{t'}$ if and only if $t' \in h$.[4] We use A^h to denote the actions available at histories in h, and h^a is the (unique) information set such that $a \in A^{h^a}$.
   (a) **Recall:** If $(a_j)_{j=1}^{J} \in h$ then $(a_j)_{j=1}^{J'} \notin h$ for any $J' < J$.
3. **Players:** A finite set of *players* $\{0\} \cup N$ and a function ι that assigns to each information set to a member of $\{0\} \cup N$.
   (a) **Nature:** Player 0 is called *nature*, its information sets $H_0 \equiv \iota^{-1}(0)$ are singletons, and for each $a \in h \in H_0$, it plays a with probability $\pi(a) > 0$.
   (b) **Payoffs:** Each player $i \in N$ has an associated *payoff function* $u_i : Z \to \mathbb{R}$.

A *(behavior) strategy profile* σ is a map from each information set $h \in H$ to a distribution over its actions, $\sigma^h \in \Delta(A^h)$, such that $\sigma^h = \pi^h$ whenever $h \in H_0$. We let Σ be the set of strategy profiles. Given that each action $a \in A$ only belongs to one information set h^a, it is convenient to use $\sigma(a)$ to denote $\sigma^{h^a}(a)$.

[4]We assume that each action only belongs to a unique information set.

3 Likelihood Assignments

In this section we define two important concepts in this paper: likelihood assignments and the assignments they generate. In the rest of the paper, an extensive-form game is assumed to be fixed.

Definition 3.1. A *likelihood assignment* is a pair (λ, f), with $\lambda : A \to \mathbb{R}_+$ and $f : A \to \mathbb{R}_{++}$, such that, for each information set $h \in H$, the set $A^h_\lambda \equiv \{a \in A^h | \lambda(a) = 0\}$ is not empty, $\sum_{a\in A^h_\lambda} f(a) = 1$ for all $h \in H$ and $f(a) = \pi(a)$ whenever $h^a \in H_0$.

As it will become evident in the paper, for each action a, the natural interpretation of $\lambda(a)$ is the limit "(likelihood) order" with which it is played in the limit of some sequence of strategy profiles of the game. If it is of order 0, $\lambda(a) = 0$, it means that it is played with positive probability, conditional on the information set h^a being reached. We also interpret $f(a)$ as the limit "(likelihood) intensity" with which an action a is played which, for actions of order 0, is just their probability. The connection of likelihood assignments to LPSs and CPSs in Appendix B provides additional intuition to these interpretations.

For each history of the game $t = (a_j)_{j=1}^J \in T$, we define (with some abuse of notation) its order and intensity as

$$\lambda(t) \equiv \textstyle\sum_{j=1}^J \lambda(a_j) \text{ and } f(t) \equiv \prod_{j=1}^J f(a_j)\ .$$

So, the order of a history is simply computed by adding up the orders of the actions that lead to it, while its intensity is the product of their intensities. It makes intuitive sense: the more likely are the actions that lead to a particular history (i.e., the lower their orders are, or the higher their intensities are), the more likely is the history itself. We can similarly assign (again with some abuse of notation) an order to each information set $h \in H$ as the lowest order of the histories that compose it, $\lambda(h) \equiv \min\{\lambda(t)|t \in h\}$, and we use h_λ to denote the set of such histories, $h_\lambda \equiv \{t \in h|\lambda(t) = \lambda(h)\}$.

An *assessment* (σ, μ) is composed by a strategy profile $\sigma \in \Sigma$ and system of beliefs μ, which assigns to each information set h a probability distribution over the histories it contains, $\mu^h \in \Delta(h)$. Using the previous definitions, we can associate an assessment to a likelihood assignment as follows:

Definition 3.2. (σ, μ) is *generated* by the likelihood assignment (λ, f) if:

1. For each $h\in H$, we have $\sigma(a) = f(a)$ for all $a \in A^h_\lambda$, and $\sigma(a) = 0$ otherwise.
2. For each $h\in H$, we have $\text{supp}(\mu^h)=h_\lambda$ and $\mu^h(t)=\frac{f(t)}{\sum_{t'\in h} f(t')}$ for all $t\in h_\lambda$.

It follows from the previous definitions that a likelihood assignment always generates a unique assessment. Also, notice that the concept of assessment generated by a likelihood assignment is consistent with our previous interpretation of the components of a likelihood assignment as a likelihood order and likelihood intensity. Indeed, in a given information set, the support of the beliefs is equal to the set of histories contained in the information set with the lowest order, since histories of higher order are interpreted to be of negligible relative likelihood. Also, the relative weight that histories of the same order receive is given by f, highlighting the interpretation of f as the intensity with which each action is played within each order.

3.1 Consistent Assessments

An assessment (μ, σ) is said to be *consistent* if there is a sequence $(\sigma_n, \mu_n)_n$ of fully-mixed assessments converging to (μ, σ) such that $\mu_n^h(t) = \Pr(t|\sigma_n)/\Pr(h|\sigma_n)$ for all $h \in H$, $t \in h$ and n.

In applications, finding a sequence of fully-mixed assessments that converges to a given assessment is not trivial, and it may be even more difficult proving that an assessment is not consistent by showing that no such a sequence exists. The following result provides a characterization of consistent assessments using likelihood assignments . It does not involve working with sequences of assessments. Its proof is constructive: for each likelihood assignment generating an assessment finds a sequence supporting it, and for each sequence supporting an consistent assessment finds a likelihood assignment generating it. Also, Example 3.1 below illustrates the need of the conditions in Definition 3.1 to make the assessment generated by a likelihood assignment consistent.

Theorem 3.1. *An assessment is consistent if and only if it is generated by a likelihood assignment.*

Proof. **"Only if" part:** Fix a consistent assessment (σ, μ) and a sequence of fully mixed strategies $(\sigma_n)_n$ with corresponding beliefs $(\mu_n)_n$ (obtained using Bayes' rule) converging to it. For each pair of actions $a, a' \in A$, define

$$p(a, a') \equiv \lim_{n \to \infty} \frac{\sigma_n(a)}{\sigma_n(a) + \sigma_n(a')}$$

and assume, without loss of generality, that it is well defined (notice that, for each n, the argument of the limit function belongs to $(0, 1)$). Notice that $p(a, a') = 1 - p(a', a)$.

We want to find some likelihood assignment (λ, f) generating (σ, μ). To do this, we will first define a function $\kappa : A \to \{0\} \cup \mathbb{N}$ that will categorize the actions in A in terms of their likelihood of being played. We will then obtain that λ can be defined so $\lambda(a) \geq \lambda(a')$ if and only if $\kappa(a) \geq \kappa(a')$, for all $a, a' \in A$. We will finally obtain f.

1. **Obtaining κ:** Define the relation $\succeq$ in A such that, for all $a, a' \in A$, $a \succeq a'$ if and only if $p(a, a') > 0$. The relation $\succeq$ is a complete preorder, since it is clearly

complete and reflexive and, if $p(a,a')>0$ and $p(a',a'')>0$, then

$$\frac{\sigma_n(a'')}{\sigma_n(a)} = \frac{\sigma_n(a')}{\sigma_n(a)}\frac{\sigma_n(a'')}{\sigma_n(a')} \xrightarrow[n\to\infty]{} \frac{1-p(a,a')}{p(a,a')}\frac{1-p(a',a'')}{p(a',a'')} \in [0,\infty) \ .$$

Therefore, $\succeq$ generates a partition $\mathcal{A}$ of the action set A and a complete order on it, denoted with some abuse of notation $\succeq$. Let A^0 be the set of actions $a\in A$ such that $\sigma(a)>0$, which is the highest set in the partition. For all $a\in A^0$, define $\kappa(a)\equiv 0$. Recursively, for each fixed $k\geq 1$ until $\cup_{k'=0}^{k-1}A^{k'}=A$, let A^k be the k-st element of the partition $\mathcal{A}$ in the $\succeq$ order, and let $\kappa(a)\equiv k$ for all $a\in A^k$. We use K to denote the value of k where the recursion stops, that is, when $\cup_{k'=0}^{K}A^{k'}=A$. We also, for the rest of the proof, fix some $\bar{a}^k\in A^k$ for each $k\in\{0,...,K\}$.

Notice that, for all $a\in A^k$, $p(a,a')\in(0,1)$ if and only if $a'\in A^k$, that is, if $\kappa(a)=\kappa(a')$. Indeed, since for all $a,a'\in A$ we have that $p(a,a')\in[0,1]$ and $p(a,a')=1-p(a',a)$, we have that $a,a'\in A^k$ if and only if $a\succeq a'\succeq a$, that is, if $0<p(a,a')=1-p(a',a)<1$. Conversely, if $\kappa(a)<\kappa(a')$, we have $p(a,a')=1$, and if $\kappa(a)>\kappa(a')$, then $p(a,a')=0$.

2. **Finding λ:** Define, for each $k\in\{0,...,K\}$, $q_n^k\equiv\sigma_n(\bar{a}^k)$. We define, for each k, $\alpha^k\in\mathbb{R}^K$ and $G^k\subset\{0,...,K\}$ in a recursive manner as follows. For $k=0$ choose $\alpha^0\equiv 0\in\mathbb{R}^K$ and $G^0=\emptyset$. Recursively, for each $k>0$, $\alpha^k\in\mathbb{R}^K$ is chosen to be, if it exists, the unique vector such that $\alpha^k_{k'}=0$ for all $k'\notin G^{k-1}$ and satisfies

$$\lim_{n\to\infty} q_n^k \exp(q_n,\alpha^k)\in(0,\infty) \ , \quad \text{where } \exp(q_n,\alpha^k)\equiv\prod_{k'=1}^{K}(q_n^{k'})^{\alpha^k_{k'}} \ , \tag{3.1}$$

in which case $G^k\equiv G^{k-1}$. If no such vector α^k exists, then assign $\alpha^k\in\mathbb{R}^K$ to be such that $\alpha^k_k=1$ and $\alpha^k_{k'}=0$ for all $k'\neq k$, and $G^k=G^{k-1}\cup\{k\}$.

Once finished (i.e., when $k=K$), define $G\equiv G^K$,[5]

$$\begin{aligned}\overline{\alpha} &\equiv \max\left\{ |\alpha^k_{k'}| \,\middle|\, k,k'\in\{0,...,K\} \right\} \ \text{ and}\\ \underline{\alpha} &\equiv \min\left\{ |\alpha^k_{k'}| \,\middle|\, k,k'\in\{0,...,K\} \text{ and } \alpha^k_{k'}\neq 0\right\} \ .\end{aligned}$$

Also, let's define the ranking function $r: G\to\{0,...,|G|-1\}$ as follows:

$$r(k)<r(k') \quad\Leftrightarrow\quad \lim_{n\to\infty}\frac{q_n^k}{q_n^{k'}}=\infty \ .$$

Notice that r is well defined and bijective. Assign, to each $k\in\{0,...,K\}$, order $\lambda(k)\equiv\sum_{k'\in G}\alpha^k_{k'}\left(1+\overline{\alpha}/\underline{\alpha}\right)^{r(k')}$. It is important to realize that $\lambda(k)$ is always positive, and it is 0 if and only if $k=0$. Indeed, since $\alpha^0=0$, clearly $\lambda(0)=0$.

[5]It is easy to see that, if $k,k'\in G$ and $k<k'$ then $r(k)<r(k')$.

Also, for each $k > 0$, we have $\lim_{n\to\infty} q_n^k = 0$, so if $\{k'\} = \arg\max_{\{k''\in G|\alpha_{k''}^k \neq 0\}} r(k'')$ then necessarily $\alpha_{k'}^k > 0$. Therefore, if $k > 0$,

$$\lambda(k) \geq \underline{\alpha}\left(\frac{\underline{\alpha}+\overline{\alpha}}{\underline{\alpha}}\right)^{r(k)} - \sum_{r=0}^{r(k)-1} \overline{\alpha}\left(\frac{\underline{\alpha}+\overline{\alpha}}{\underline{\alpha}}\right)^{r(k')}$$
$$> \underline{\alpha}\left(\frac{\underline{\alpha}+\overline{\alpha}}{\underline{\alpha}}\right)^{r(k)} - \sum_{r=-\infty}^{r(k)-1} \overline{\alpha}\left(\frac{\underline{\alpha}+\overline{\alpha}}{\underline{\alpha}}\right)^{r(k')} = 0 \ .$$

Define finally, for each $a \in A$, $\alpha(a) \equiv \alpha^{\kappa(a)}$ and, with some abuse of notation, $\lambda(a) \equiv \lambda(\kappa(a))$.

3. **λ applied to histories:** Define, for each $t = (a_j)_{j=1}^J \in T$, $\alpha(t) \equiv \sum_{j=1}^J \alpha(a_j)$ and $\lambda(t) \equiv \sum_{j=1}^J \lambda(a_j)$, so notice that $\lambda(t) \equiv \sum_{k'\in G} \alpha(t)_{k'}^k \left(1 + \bar{\alpha}/\underline{\alpha}\right)^{r(k')}$. Take two different histories $t, t' \in T$, with $t = (a_j)_{j=1}^J$ and $t = (a_j')_{j=1}^{J'}$. Then,

$$\frac{\Pr(t|\sigma_n)}{\Pr(t'|\sigma_n)} = \frac{\prod_{j=1}^J \sigma_n(a_j)}{\prod_{j=1}^{J'} \sigma_n(a_j')} = \underbrace{\frac{\prod_{j=1}^J \sigma_n(a_j)/\sigma_n(\bar{a}^{\kappa(a_j)})}{\prod_{j=1}^{J'} \sigma_n(a_j')/\sigma_n(\bar{a}^{\kappa(a_j')})}}_{\equiv(*)} \underbrace{\frac{\prod_{j=1}^J \sigma_n(\bar{a}^{\kappa(a_j)})}{\prod_{j=1}^{J'} \sigma_n(\bar{a}^{\kappa(a_j')})}}_{\equiv(**)} \ .$$

As $n \to \infty$, the term $(*)$ converges to some value in $(0, \infty)$. Indeed, for each $a \in A$, by the definition of $\bar{a}^{\kappa(a)}$, we have that $\lim_{n\to\infty} \sigma_n(a_j)/\sigma_n(\bar{a}^{\kappa(a_j)}) \in (0, 1)$. Using the definition of α, we have the following expression for the limit of the term $(**)$:

$$\lim_{n\to\infty} (**) \propto \lim_{n\to\infty} \frac{\exp\left(q_n, \sum_{j=1}^J \alpha(a_j)\right)}{\exp\left(q_n, \sum_{j=1}^{J'} \alpha(a_j')\right)} = \lim_{n\to\infty} \frac{\exp\left(q_n, \sum_{j=1}^J \alpha(t)\right)}{\exp\left(q_n, \sum_{j=1}^{J'} \alpha(t')\right)}$$
$$= \lim_{n\to\infty} \exp(q_n, \alpha(t) - \alpha(t')) \ .$$

As a result, $\lim_{n\to\infty} \frac{\Pr(t|\sigma_n)}{\Pr(t'|\sigma_n)} \in (0, 1)$ if and only if $\alpha(t) = \alpha(t')$, so $\lambda(t) = \lambda(t')$, and $\lim \frac{\Pr(t|\sigma_n)}{\Pr(t'|\sigma_n)} = 0$ if and only if $\lambda(t) > \lambda(t')$. This implies that, for all $h \in H$, we have that $\mathrm{supp}(\mu^h) = h_\lambda$.

4. **Finding f:** For each $k \in G$, pick some $\xi^k \in \mathbb{R}_{++}$, and pick $\xi^k \equiv 1$ for all $k \notin G$.[6] Define, for each $a \in A$,

$$f_n(a) \equiv \frac{\exp(\xi, \alpha(a))}{\exp(q_n, \alpha(a))} \sigma_n(a) \ . \tag{3.2}$$

Notice that $f(a) \equiv \lim_{n\to\infty} f_n(a)$ is a well defined, positive number, and also that if $\lambda(a) = 0$ (so $\alpha = 0$), we have $f_n(a) = \sigma_n(a)$ (since, in this case, $\alpha(a) = 0$, so

[6]For this proof, one can take $\xi^k = 1$ for all k, but it is convenient to define f_n in a more general way for arguments in other proofs of this paper.

both the numerator and the denominator of the fraction on the right hand side of equation (3.2) are 1.) Define also, for each $t=(a_j)_{j=1}^J \in T$, $f(a) \equiv \prod_{j=1}^J f(a_j)$. So, we have

$$\Pr(t|\sigma_n) = \prod_{j=1}^J \sigma_n(a_j) = \frac{\exp(q_n, \alpha(t))}{\exp(\xi, \alpha(t))} \prod_{j=1}^J f_n(a_j) \ .$$

Then, since $\alpha(t)$ is the same for all $t \in h_0$, the following holds:

$$\mu_n^h(t) = \frac{\frac{\exp(q_n,\alpha(t))}{\exp(\xi,\alpha(t))} f_n(a)}{\sum_{t' \in h_\lambda} \frac{\exp(q_n,\alpha(t))}{\exp(\xi,\alpha(t))} \prod_{a \in t} f_n(a)} \xrightarrow[n\to\infty]{} \frac{f(t)}{\sum_{t' \in h_\lambda} f(t')} \ ,$$

and for $t \notin h_\lambda$, we have that $\lim_{n\to\infty} \mu_n^h(t) \to 0$. This concludes the first part of the proof.

"If" part: Assume that (σ, μ) is an assessment generated by a likelihood assignment (λ, f). We aim at defining a sequence $(\sigma_n)_n$ of fully mixed strategies such that, the corresponding sequence of completely mixed assessments $(\mu_n, \sigma_n)_n$ converges to (μ, σ). We do this by fixing a strictly decreasing sequence $(\varepsilon_n)_n$ converging to 0, and proceeding as follows:

1. Define $\sigma_n(a) = C_n^{h^a} \varepsilon_n^{\lambda(a)} f(a)$, where $C_n^{h^a}$ is a constant to make σ_n^t a probability distribution over A^{h^a}. Notice that, defining $\bar{f} = \max\{f(a)|a \in A\}$ we have

$$1 \geq C_n^h \geq 1 - \sum_{\ell \in \lambda(A)} \varepsilon_n^\ell \, |A| \, \bar{f} \ .$$

   So, clearly, as $n \to \infty$, $C_n^h \to 1$. As a result, it is clear that $\sigma_n \to \sigma$.

2. Now, notice that, for all $t = (a_j)_{j=1}^J$, we have

$$\Pr(t|\sigma_n) = \prod_{j=1}^J \sigma_n(a_j) = \varepsilon_n^{\sum_{j=1}^J \lambda(a_j)} \prod_{j=1}^J C_n^{h_{a_j}} f(a) = \varepsilon_n^{\sum_{j=1}^J \lambda(a_j)} f(t) \prod_{j=1}^J C_n^{h_{a_j}} \ .$$

   Hence, as n increases, μ_n^h converges to some μ_∞^h with support equal to h_λ and, for each $t = (a_j)_{j=1}^J \in h_\lambda$, we have

$$\mu_\infty^h(t) = \lim_{n\to\infty} \frac{\Pr(t|\sigma_n)}{\Pr(h_\lambda|\sigma_n)} = \lim_{n\to\infty} \frac{f(t) \prod_{j=1}^J C_n^{h_{a_j}}}{\sum_{(a'_j)_{j=1}^{J'} \in h_\lambda} f((a'_j)_{j=1}^{J'}) \prod_{j=1}^{J'} C_n^{h_{a_j}}} = \frac{f(t)}{\sum_{t' \in h_\lambda} f(t')} \ .$$

   So, $\mu_\infty = \mu$ and the proof is concluded.

$\square$

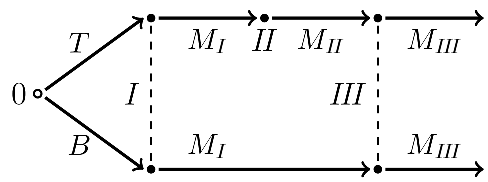

Figure 1: Game from Example 3.1.

Example 3.1. This example illustrates necessity of the conditions that λ and f in a likelihood assignment (Definition 3.1) to ensure the consistency of the assessment it generates (using their additive and multiplicative properties over histories). See Figure 1. In the game, first nature decides T or B, assume with the same probability $\frac{1}{2}$. Then, necessarily, player I assigns probability $\frac{1}{2}$ to both the bottom and top branches. Also, clearly, in any consistent assessment, player III does the same. Then, in order for the assessment generated by a likelihood assignment in the game (Definition 3.2) to be consistent, it must be that $\lambda(M_{II}) = 0$ and $f(M_{II}) = 1$.

4 Stable Outcomes

4.1 Optimality

In an extensive form game, defining the continuation value of a player i at some information set $h \in H_i$ for a given assessment (μ, σ) requires defining the continuation outcome after each possible action choice in A^h. Since doing so is notationally intensive, we abstain from it in this paper (see Osborne and Rubinstein (1994), Section 12.2, for details). Instead, we provide a characterization of it using likelihood assignments.

Fix a likelihood assignment (λ, f). For all $a \in A$, we then define

$$u^{\lambda,f}(a) \equiv \sum_{z \in Z^a_\lambda} \frac{f(z)}{f(a)} u_{\iota(h^a)}(z) \ ,$$

where Z^a_λ is the set of terminal histories $z \in Z$ that have a as one of their elements and such that $\lambda(z) - \lambda(a)$ is minimal. We interpret $\lambda(z) - \lambda(a)$, and $f(z)/f(a)$ as, respectively, the order and the intensity of the terminal history z if player $\iota(h^a)$ plays a for sure if the information set h^a is reached, while the random choices in the other information sets (hers and the other players') remain the same.

We use $a \in \mathrm{PBR}(\lambda, f)$ to denote $u_i^{\lambda,f}(a) \geq u_i^{\lambda,f}(a')$ for all $a' \in A^h$. The reader may verify that, dividing $u^{\lambda,f}(a)$ by $\sum_{z \in Z^a_\lambda} f(z)/f(a)$, one obtains the continuation payoff (in the usual sense) at information set h^a after playing a (for the assessment generated by (λ, f)), and that $a \in \mathrm{PBR}(\lambda, f)$ if and only if a maximizes the continuation payoff.

4.2 Perturbations

A *behavioral tremble* is a sequence of functions $(\eta_n : A \to \mathbb{R}_{++})_n$ such that, for each $a \in A$, $\lim_{n\to\infty} \eta_n(a)$ converges to 0 and, for each $h \in H$ and $n \in \mathbb{N}$, we have $\sum_{a\in h} \eta_n(a) < 1$. The n-th perturbed game, denoted $G(\eta_n)$, is such that, for each action $a \in A$, player $\iota(a)$ is required to choose a with probability at least $\eta_n(a)$, that is, $\sigma(a) \geq \eta_n(a)$.

We now introduce the concept of perturbed likelihood assignment, which is analogous to behavioral trembles.

Definition 4.1. A *perturbed likelihood assignment* is a pair $(\tilde{\lambda}, \tilde{f})$, with $\tilde{\lambda}, \tilde{f} : A \to \mathbb{R}_{++}$.

In the same way that perturbations are not sequences of strategy profiles, the definition of perturbed likelihood assignments relaxes the conditions that a likelihood assignment satisfies: a perturbed likelihood assignment differs from a likelihood assignment in that the likelihood orders of all actions are strictly positive, and therefore there is no condition that $\tilde{f}$ satisfies (except for being strictly positive). The proof of Theorem 4.1 shows that we can associate a behavioral tremble to a perturbed likelihood assignment and, using an argument similar to the one used in the proof of Theorem 3.1, one can see that we can also associate perturbed likelihood assignments to behavioral trembles.

Definition 4.2. Fix a likelihood assignment (λ, f) and a perturbed likelihood assignment $(\tilde{\lambda}, \tilde{f})$. Then, (λ, f) is an *assignment equilibrium* for $(\tilde{\lambda}, \tilde{f})$ if, for all $a \in A$,

1. **Superseding:** $\lambda(a) \leq \tilde{\lambda}(a)$, and $f(a) \geq \tilde{f}(a)$ if $\lambda(a) = \tilde{\lambda}(a)$.
2. **Optimality:** $a \in \text{PBR}(\lambda, f)$ if either $\lambda(a) < \tilde{\lambda}(a)$, or $\lambda(a) = \tilde{\lambda}(a)$ and $f(a) > \tilde{f}(a)$.

Our definition of "superseding" sheds light on the connection between the concepts of perturbed likelihood assignment and behavioral tremble: conditional on a given information set being reached, the likelihood that the corresponding player can assign to an action has to be no lower than the likelihood with which the perturbation generates it. So, for a fixed a perturbation $(\tilde{\lambda}, \tilde{f})$, an assessment (λ, f) supersedes $(\tilde{\lambda}, \tilde{f})$ only if, at each information set, it assigns a weakly higher "likelihood" to each of the actions in such information set. Furthermore, if an assignment equilibrium for $(\tilde{\lambda}, \tilde{f})$ assigns a strictly higher likelihood than $(\tilde{\lambda}, \tilde{f})$ to an action, such an action is a best response.

In order to illustrate the concept of assignment equilibrium, we present now a lemma characterizing sequential equilibria. This result is also going to be helpful to get shed some light on the next result about stable outcomes.

Lemma 4.1. *(μ, σ) is a sequential equilibrium if and only if it is generated by some assignment equilibrium (for some perturbed likelihood assignment $(\tilde{\lambda}, \tilde{f})$.)*

Proof. **"Only if" part:** Let (μ, σ) be sequential equilibrium, and let (λ, f) be a likelihood assignment generating it (which exists by Theorem 3.1). Take, the perturbed likelihood assignment

$$(\tilde{\lambda}(a), \tilde{f}(a)) \equiv \begin{cases} (1,1) & \text{if } \lambda(a) = 0, \\ (\lambda(a), f(a)) & \text{if } \lambda(a) > 0. \end{cases}$$

Clearly, $(\tilde{\lambda}, \tilde{f})$ is superseded by (λ, f). Also, notice that given that $\lambda(a) < \tilde{\lambda}(a)$ if and only if $\sigma(a) > 0$, so a is optimal (because of optimality in sequential equilibria). Finally, for all a such that $\sigma(a) = 0$ we have $\tilde{\lambda}(a) = \lambda(a)$ and $\tilde{f}(a) = f(a)$, and therefore (λ, f) optimally supersedes $(\tilde{\lambda}, \tilde{f})$.

"If" part: Let (λ, f) be a likelihood assignment optimally superseding some perturbed likelihood assignment $(\tilde{\lambda}, \tilde{f})$, and let (μ, σ) be the assessment generated by (λ, f). Given that (λ, f) is consistent (by Theorem 3.1), it is only left to verify that $a \in \mathrm{PBR}(\mu, \sigma)$ for all a such that $\sigma(a) > 0$. This follows trivially from the fact that whenever a is such that $\mu(a) > 0$ we have $\lambda(a) = 0 < \tilde{\lambda}(a)$ and therefore, by the fact that (λ, f) optimally supersedes $(\tilde{\lambda}, \tilde{f})$, we have $a \in \mathrm{PBR}(\mu, \sigma)$. □

4.3 Stable Outcomes

An *outcome* o is a distribution over the terminal histories Z. A strategy profile σ *leads* to some outcome o if the distribution over terminal histories induced by σ is o. We say that (λ, f) leads to o if the strategy profile in the assessment generated by (λ, f) leads to o. Notice that o it can be computed as

$$o(z) = \begin{cases} f(z)/(\sum_{z' \in Z_\lambda} f(z')) & \text{if } z \in Z_\lambda, \\ 0 & \text{if } z \notin Z_\lambda, \end{cases}$$

where, consistently with the previous notation, Z_λ is the set of terminal histories $z \in Z$ with $\lambda(z) = 0$.

Kreps and Wilson (1982) show that a generic extensive form game has a finite number of (Nash-)equilibrium outcomes. Kohlberg and Mertens (1986) show that, generically, there is at least one of these outcomes o which is stable: if the game is perturbed (through a "small" tremble), the resulting game has an equilibrium outcome which is close to o. So, in a generic game, an outcome o is stable only if for any behavioral tremble $(\eta_n)_n$ there is a sequence $(\sigma_n)_n$, where σ_n is a Nash equilibrium of $G(\eta_n)$, such that the limit of the outcomes of σ_n is o.

Even though strategic stability is a sensible requirement to the behavior expected by strategic agents (stable sets of equilibria satisfy invariance, admissibility, and iterated dominance, see Kohlberg and Mertens (1986) for the corresponding definitions,) it is, in general, difficult to prove that an outcome is stable. So, in a similar spirit, we define assignment-stable outcomes using likelihood assignments:

Definition 4.3. An outcome o is *assignment-stable* if, for each perturbed likelihood assignment $(\tilde{\lambda}, \tilde{f})$, there is an assignment equilibrium for $(\tilde{\lambda}, \tilde{f})$ leading to o.

We now establish that a stable outcomes are also assignment-stable. This has two implications. The first is that assignment-stable outcomes exist in generic extensive-form games. Second, when a game has a unique assignment-stable outcome, then it is also stable. So, since for some games it easy to show that they have a unique assignment-stable outcome, such outcome is also stable. As the examples below show, this result is useful to rule out the stability of particular outcomes of particular dynamic games, as one may have an intuition about what perturbations may "destroy" the outcome.

Theorem 4.1. *An outcome is stable only if it is assignment-stable.*

Proof. Assume that the outcome o is stable. Fix a perturbed likelihood assignment $(\tilde{\lambda}, \tilde{f})$. For a strictly decreasing sequence $(\varepsilon_n)_n$ converging to 0, define the following behavioral tremble $\eta_n(a) = \varepsilon_n^{\tilde{\lambda}(a)} \tilde{f}(a)$, where ε_1 is taken to be small enough that, for each $h \in H$, we have $\sum_{a\in A^h} \eta_n(a) < 1$. Define $\eta_n(h) \equiv \sum_{a\in A^h} \eta_n(a)$.

Let $(\mu_n, \sigma_n)_n$ be a sequence of sequential equilibria of the perturbed game $G(\eta_n)$ with outcome converging to outcome o. Define $\hat{\sigma}_n(a) \equiv \frac{\sigma_n(a)-\eta_n(a)}{1-\eta_n(h_a)}$, and note that it is a strategy profile in Σ. Let μ_n be the belief system generated by σ_n (note that σ_n has full support). Assume, without loss of generality, that $(\hat{\sigma}_n)_n$, $(\sigma_n, \mu_n)_n$ converge, and let $\hat{\sigma}$ and (σ, μ) be the corresponding limits. For each $a \in A$ define

$$\gamma(a) \equiv \lim_{n\to\infty} \frac{\eta_n(a)}{\sigma_n(a)} .$$

Notice that, since $0 < \eta_n(a) \leq \sigma_n(a)$, the Bolzano-Weierstrass ensures that there is a subsequence of the original sequence where $\gamma(a)$ exists for all $a \in A$. To simplify notation, assume that this is the case for the main sequence.

Let (λ^*, f^*) be the likelihood assignment associated to the (consistent) assessment (σ, μ), obtained as in the "Only if" part of the proof of Theorem 3.1. Let K, $\alpha \in \mathbb{R}^K \times \mathbb{R}^K$ and $G \subset \{0, ..., K\}$ be the results of the algorithm.

Assume first that $\gamma(a) = 0$ for all $a \in A$. This implies that, for all $a \in A$, $\sigma_n(a) > \tilde{\sigma}_n(a)$ for n big enough, that is, all actions are optimal. As a result, for n big enough, $\hat{\sigma}_n$ has full support, and has μ as its limit beliefs. In this case, define $\bar{\lambda}^* = \max_{a\in A} \lambda^*(a)$ and $\underline{\tilde{\lambda}} \equiv \min_{a\in A} \lambda(a)$. Then, define $\lambda(a) \equiv \frac{\underline{\tilde{\lambda}}}{2\,\bar{\lambda}^*} \lambda^*(a)$ and $f \equiv f^*$. Notice that (λ, f) has (σ, μ) as its associated assessment. Furthermore, it is clear that $\lambda(a) < \tilde{\lambda}(a)$ for all $a \in A$, and that (by continuity in the payoffs) (λ, f) optimally supersedes $(\tilde{\lambda}, \tilde{f})$.

Assume then that $\gamma(a) > 0$ for some $a \in A$ (notice that this implies that $\lambda^*(a) > 0$). Let's define $k^* \equiv \kappa(a)$. We then claim that for all $a, a' \in A$ such that $\gamma(a), \gamma(a') \in (0,1]$ we have $\lambda^*(a)/\tilde{\lambda}(a) = \lambda^*(a')/\tilde{\lambda}(a)$. Indeed, notice first that

$$\lim_{n\to\infty} \frac{\sigma_n(a)^{1/\tilde{\lambda}(a)}}{\sigma_n(a)^{1/\tilde{\lambda}(a')}} = \frac{\gamma(a)^{-1/\tilde{\lambda}(a)}}{\gamma(a')^{-1/\tilde{\lambda}(a)}} \lim_{n\to\infty} \frac{\eta_n(a)^{1/\tilde{\lambda}(a)}}{\eta_n(a)^{1/\tilde{\lambda}(a')}} = \frac{\gamma(a)^{-1/\tilde{\lambda}(a)}}{\gamma(a')^{-1/\tilde{\lambda}(a)}} \frac{f(a)^{1/\tilde{\lambda}(a)}}{f(a)^{1/\tilde{\lambda}(a')}} \in (0,1) .$$

So, by the construction of α, $\alpha(a)/\tilde{\lambda}(a) = \alpha(a')/\tilde{\lambda}(a')$, and therefore $\lambda^*(a)/\tilde{\lambda}(a) = \lambda^*(a')/\tilde{\lambda}(a')$. Define then $\lambda(a) \equiv \frac{\tilde{\lambda}(\tilde{a})}{\lambda^*(\tilde{a})} \lambda^*(a)$, where $\tilde{a}$ is any action such that $\gamma(\tilde{a}) > 0$ (notice that the particular choice of $\tilde{a}$ is irrelevant). Proceeding similarly, it is easy to show that $\gamma(a) = 0$ then $\lambda(a) < \tilde{\lambda}(a)$, if $\gamma(a) \in (0,1)$ then $\lambda(a) = \tilde{\lambda}(a)$, and if $\gamma(a) = 1$ then $\lambda(a) > \tilde{\lambda}(a)$.

Let $\tilde{G} \subset G$ be the set of $k \in G$ such that there is some $a \in A^k$ such that $\gamma(a) > 0$, and let $\tilde{\lambda}(k)$ be defined as $\tilde{\lambda}(a)$ (note that if there is another action $a' \in A^k$ such that $\gamma(a') > 0$, then necessarily $\tilde{\lambda}(a') = k$). Now, recall that in the definition of f^* (from equation (3.2)) there is freedom from choosing the parameters $\xi \in \mathbb{R}^G_{++}$. Therefore, we choose f such that $\xi^k = 1$ for all $k \in G \backslash \tilde{G}$ and

$$\xi^k = \lim_{n\to\infty} \sigma_n(\bar{a}^k)/\varepsilon_n^{\tilde{\lambda}(k)} \quad \text{for all } k \in \tilde{G},$$

where $\bar{a}^k$ is defined in the proof of Theorem 3.1. Then, notice that for all $a \in A$ such that $\gamma(a) > 0$, we have

$$f(a) = \lim_{n\to\infty} \frac{\exp(\xi, \alpha(a))}{\exp(q_n, \alpha(a))} \sigma_n(a) = \lim_{n\to\infty} \frac{\sigma_n(a)}{\varepsilon_n^{\tilde{\lambda}(\kappa(a))}} = \tilde{f}(a) + \underbrace{\lim_{n\to\infty} \frac{\hat{\sigma}_n(a)}{\varepsilon_n^{\tilde{\lambda}(\kappa(a))}}}_{(*)} .$$

If $\gamma(a) = 1$ then $(*)$ is equal to 0, so $f(a) = \tilde{f}(a)$. If, instead, $\gamma(a) \in (0,1)$, then $(*)$ is strictly positive. In this case, $\hat{\sigma}_n(a) > 0$ for all n high enough. Notice finally that, if $\lambda(a) < \tilde{\lambda}(a)$ then $\gamma(a) = 0$, so we have that, if n is high enough, $\hat{\sigma}_n(a) > 0$. Therefore, by continuity in the payoffs, if $a \in A$ such that $\lambda(a) < \tilde{\lambda}(a)$, or $\lambda(a) = \tilde{\lambda}(a)$ and $f(a) > \tilde{f}(a)$, it is optimal to play a at h_a, so (λ, f) optimally supersedes $(\tilde{\lambda}, \tilde{f})$. □

4.4 Examples

Example 4.1. Figure 2 (which is borrowed from Osborne and Rubinstein (1994), Figure 227.1) provides an example of two games with the same normal form, but with different sets of sequential equilibria. Indeed, the strategy profile (B_1, B_2) is part of a sequential equilibrium on the left game, but not on the right. This indicates that the outcome assigning probability one to history B_1 is not stable. To see this, consider a perturbation likelihood assignment where $\tilde{\lambda}(T_1) < \tilde{\lambda}(M_1)$. If a likelihood assignment optimally supersedes it, (1) either $\lambda(T_1) < \lambda(M_1)$, but then player 2 chooses T_2, so player 1 has

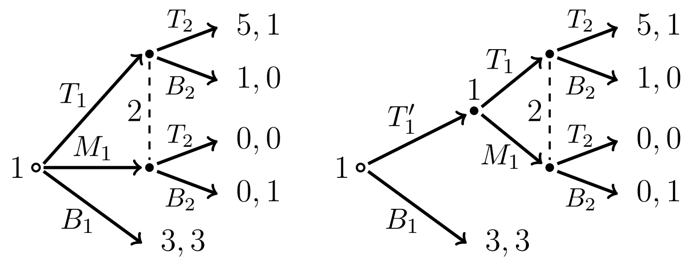

Figure 2: Game from Example 4.1.

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

Table 1: Table of example 4.2.

an incentive to deviate, (2) or $\lambda(T_1) \geq \lambda(M_1)$, but then playing M_1 should be optimal (i.e., provide a payoff of at least 3), which clearly is not.

Example 4.2 (Dynamic Signaling). 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 that is of one of two types, $\theta = L, H$. The prior about $\theta = H$ is 1/2. Only the student knows her type. The cost of education is $0 < c_H < c_L < 1$. Each firm has 3 options: make a high offer (h), make a low offer (ℓ) or make no offer ($\emptyset$). 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, and otherwise moves to the second period. The payoffs (without education costs) are given in Table 1 (the payoff can be perturbed to make them generic).[7] Notice that the expected payoff for the firm 1 if all types of the student accept for sure a low offer is 0.5, while if it offers h its expected payoff is -0.5. Therefore, in a one-period model, the (first) firm offers ℓ and the student accepts the offer for sure. We consider several cases, and in all of them $\tilde{\lambda}$ refers to the order in a perturbation likelihood assignment, while λ refers to the order in likelihood assignment which gives rise to a stable outcome.[8]

1. *Spence (static) model:* Let us first assume that, to get an offer in the first period, the student has to irreversibly exit education (i.e., go on the market), denoted E,

[7] If, for example, the type is L, the firm 1 offers ℓ and the student accepts, the payoff for the student is 3, the payoff of the firm is -3, and the payoff of the second firm is 0. If, instead, the student rejects the first offer and the second firm offers her h, the payoff of firm 1 is 0, the payoff of the student is $5 - c_L$, and the payoff of the second firm is -5.

[8] We study in a unified framework versions of the standard static signaling model (Spence, 1973) and the dynamic signaling models with public (Noldeke and Van Damme, 1990) and (Swinkels, 1999) preemptive offers.

while if she stays, denoted S, she pays the education cost and goes on the market in the second period. Assume first that there is an assignment-stable outcome that exhibits exit in the first period. If $\tilde{\lambda}(S; L) > \tilde{\lambda}(S; H)$ then, necessarily, the L-student should willing to stay in education, since otherwise firm 2 offers h, while firm 2 offers either ℓ or $\emptyset$. Nevertheless, in this case, the H-student is strictly willing to stay in education (since her education cost is low and the payoff from accepting h is high) and, as a result, firm 1 makes no offer, while firm 2 makes either offer ℓ or offer h (since its posterior is higher than 1/2), which gives the L-student a strict incentive to stay.

Hence, in the unique stable outcome, all types of student exit in the second period, and get offered ℓ. It is important to notice that, in this case, the (unique) fully stable set of equilibria contains multiple equilibria. If $\tilde{\lambda}(E, L) > \tilde{\lambda}(E, H)$, then $\lambda(E, L) = \tilde{\lambda}(E, H)$, and firm 1 randomizes between $\emptyset$ and ℓ so the L-student is indifferent on going on the market in period 0 or staying in education. If $\tilde{\lambda}(E, L) < \tilde{\lambda}(E, H)$, then $\lambda(E, L) = \tilde{\lambda}(E, L) < \tilde{\lambda}(E, H) = \lambda(E, H)$, so firm 1 offers $\emptyset$.

2. *Public offers case:* If, with positive probability on the path of play, a low offer is made by firm 1, and such an offer is accepted with positive probability, then necessarily it is accepted with positive probability by the H-student, so the L-student has a strict incentive to accept it, and hence $\lambda(R; L, \ell) = \tilde{\lambda}(R; L, \ell)$. Still, if $\tilde{\lambda}(R; H, \ell) < \tilde{\lambda}(R; L, \ell)$, then necessarily firm 2 offers h, which leads to a contradiction. So, in the unique stable outcome, firm 1 offers $\emptyset$, firm 2 offers ℓ, which is accepted by both types of student.

3. *Private offers case:* Clearly, firm 1 makes acceptable offers (equal to ℓ) with positive probability, since otherwise firm 2 would offer ℓ, but then firm 1 would benefit from offering ℓ. Assume that, on the path of play, ℓ is offered by firm 1 and accepted for sure by both types of student. Assume a perturbation likelihood assignment where $\tilde{\lambda}(R; H, \ell) < \tilde{\lambda}(R; L, \ell)$. If $\lambda(R; H, \ell) < \lambda(R; L, \ell)$, then firm 2 offers h, but then both types of the students strictly prefer to wait. If the L-student is indifferent on accepting ℓ, then the H-student strictly prefers to reject it. Then, necessarily, the H-student indifferent on accepting a low offer in the first period. As a result, in the unique stable outcome, firm 1 randomizes between $\emptyset$ and ℓ, firm 2 randomizes between ℓ and h, the L-student always accepts an offer, and the H-student accepts ℓ in the first period with a positive probability but not for sure and accepts the offer in the second period.

5 Conclusions

We provide new characterizations of consistent beliefs, sequential equilibria and stable outcomes. We do this using likelihood assignments, which are both intuitive and simple to use. Our results may help in selecting equilibria consistently across applications, and therefore may contribute to reconcile the theoretically desirable refinements and the ones used in practice.

References

Banks, Jeffrey S and Joel Sobel (1987) "Equilibrium selection in signaling games," *Econometrica: Journal of the Econometric Society*, pp. 647–661.

Battigalli, Pierpaolo (1996) "Strategic independence and perfect Bayesian equilibria," *Journal of Economic Theory*, Vol. 70, pp. 201–234.

Blume, Lawrence, Adam Brandenburger, and Eddie Dekel (1991) "Lexicographic probabilities and equilibrium refinements," *Econometrica: Journal of the Econometric Society*, pp. 81–98.

Cho, In-Koo and David M Kreps (1987) "Signaling games and stable equilibria," *The Quarterly Journal of Economics*, Vol. 102, pp. 179–221.

Fudenberg, Drew and Jean Tirole (1991) "Perfect Bayesian equilibrium and sequential equilibrium," *journal of Economic Theory*, Vol. 53, pp. 236–260.

Govindan, Srihari and Tilman Klumpp (2003) "Perfect equilibrium and lexicographic beliefs," *International Journal of Game Theory*, Vol. 31, pp. 229–243.

Kohlberg, Elon and Jean-Francois Mertens (1986) "On the strategic stability of equilibria," *Econometrica*, Vol. 54, pp. 1003–1037.

Kohlberg, Elon and Philip J Reny (1997) "Independence on relative probability spaces and consistent assessments in game trees," *Journal of Economic Theory*, Vol. 75, pp. 280–313.

Kreps, David M and Robert Wilson (1982) "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, pp. 863–894.

Mailath, George J, Larry Samuelson, and Jeroen M Swinkels (1997) "How proper is sequential equilibrium?" *Games and Economic Behavior*, Vol. 18, pp. 193–218.

Myerson, Roger B (1978) "Refinements of the Nash equilibrium concept," *International journal of game theory*, Vol. 7, pp. 73–80.

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

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

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

Osborne, Martin J and Ariel Rubinstein (1994) *A course in game theory*: MIT press.

Selten, Reinhard (1975) "Reexamination of the perfectness concept for equilibrium points in extensive games," *International journal of game theory*, Vol. 4, pp. 25–55.

Spence, Michael (1973) "Job market signaling," *The quarterly journal of Economics*, Vol. 87, pp. 355–374.

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

A Omitted Proofs

Proof of Corollary B.1. By Theorem 3.1, the "if part" is clear. To prove the "only if" part assume that an assessment (μ, σ) is consistent, and let (λ, f) generate (μ, σ) (which by Theorem 3.1 exists.) Our goal is to find an likelihood assignment $(\hat{\lambda}, \hat{f})$ such that $\hat{f} = f$ and $\hat{\lambda}(t) \geq \hat{\lambda}(t')$ if and only if $\lambda(t) \geq \lambda(t')$, and also satisfying that $\hat{\lambda}(a)$ is a rational number for all $a \in A$. Notice that, if such a likelihood assignment exists, it generates (μ, σ), so showing its existence proves our result, since once obtained one can multiply $\hat{\lambda}$ by the least common divisor of all $\{\hat{\lambda}(a) | a \in A\}$ and obtain an integer likelihood assignment.[9]

Fix some $\varepsilon > 0$ to be determined. We use the function $\lambda : \{0, ..., K\} \to \mathbb{R}_+$ defined in the proof of Theorem 3.1. For each $k \in \{0, ..., K\}$, we define $\hat{\lambda}_\varepsilon : \{0, ..., K\} \to \mathbb{Q}_+$ and Q^k recursively as follows. First, assign $\hat{\lambda}_\varepsilon(0) \equiv 0$ and $Q^0 \equiv \emptyset$. Then, for all $k > 0$ until $k = K$, proceed recursively as follows:

1. If $\lambda(k)$ is a rational combination of elements in $\{\lambda(k') | k' \in Q^{k-1}\}$, given by $\lambda(k) = \sum_{k' \in Q^{k-1}} \gamma_\varepsilon^{k,k'} \lambda(k')$ (where $\gamma_\varepsilon^{k,k'} \in \mathbb{Q}$ for all $k' \in Q^{k-1}$), then assign $Q^k \equiv Q^{k-1}$, $\hat{\lambda}_\varepsilon(k) = \sum_{k' \in Q^{k-1}} \gamma_\varepsilon^{k,k'} \hat{\lambda}_\varepsilon(k') \in \mathbb{Q}$.
2. Otherwise, if $\lambda(k)$ is a not rational combination of elements in $\{\lambda(k') | k' \in Q^{k-1}\}$, assign $\hat{\lambda}_\varepsilon(k)$ a rational number in $(\lambda(k) - \varepsilon, \lambda(k) + \varepsilon)$ and $Q^k \equiv Q^{k-1} \cup \{k\}$.

[9]It is trivial to show that if a likelihood assignment (λ, f) generates (μ, σ), then $(\xi\, \lambda, f)$ also generates (μ, σ) for all $\xi \in \mathbb{R}_{++}$.

Let Q denote Q^K. Notice that, by construction, $\hat{\alpha}_k^{k'} \in \mathbb{Q}^K$ for all $k, k' \in \{0, ..., K\}$. Also, for every $k, k' \in \{0, ..., K\}$, there exists a unique rational combination such that $\hat{\lambda}_\varepsilon(k) = \sum_{k' \in Q} \gamma_\varepsilon^{k,k'} \hat{\lambda}_\varepsilon(k')$.

Notice that if $\ell^\dagger$ is a rational combination of $\{\lambda(k)|k = 0, ..., K\}$, then there exists a unique set of rational coefficients $\{\gamma_\varepsilon^k(\ell^\dagger)|k = 0, ..., K\}$ such that $\ell^\dagger = \sum_{k \in Q} \gamma_\varepsilon^k(\ell^\dagger)\, \lambda(k)$. Indeed, by construction, $\lambda(k)$ for each $k \in Q$ cannot be written as a rational combination of $\{\lambda(k')|k' \in Q \backslash \{k\}\}$, while for all $k \notin Q$ we have that $\lambda(k)$ can be written uniquely as a rational combination of elements in $\{\lambda(k)|k \in Q\}$.

Clearly, as $\varepsilon \searrow 0$, we have that $\hat{\lambda}_\varepsilon(k) \to \lambda(k)$. So, if ε is small enough, $(\hat{\lambda}_\varepsilon(\cdot), f)$ is a likelihood assignment . Also, by continuity, if $t, t' \in T$ are such that $\lambda(t) > \lambda(t')$, we have that $\hat{\lambda}_\varepsilon(t) > \hat{\lambda}_\varepsilon(t)$ if ε is small enough, where as before, with some abuse of notation, we denote $\hat{\lambda}_\varepsilon(\cdot) \equiv \hat{\lambda}_\varepsilon(\kappa(\cdot))$. Finally, assume that $t, t' \in T$ are such that $\lambda(t) = \lambda(t')$. Let $\{\gamma_\varepsilon^k(t)|k = 0, ..., K\}$ be the unique set of rational coefficients such that $\lambda(t) = \sum_{k \in Q} \gamma_\varepsilon^k(t)\, \lambda(k)$. Then, it is clear that

$$\hat{\lambda}_\varepsilon(t) = \sum_{k \in Q} \gamma_\varepsilon^k(t)\, \hat{\lambda}_\varepsilon(k) = \hat{\lambda}_\varepsilon(t') \ .$$

This concludes the proof. $\square$

B Relationship to LPSs and CPSs

In this section, we shed light on the link between, on one hand, likelihood assignments, and on the other hand lexicographic probability systems (LPSs) and conditional probability systems (CPSs), as they have also been used to characterize consistent assessments.

B.1 Integer Likelihood Assignments

The proof of Theorem 3.1 obtains an integer measure, κ, of the limit order of each action for a given sequence of fully-mixed assessments. Still, this measure lacks the additive properties of λ, that is, the order of a history cannot be obtained by simply adding the order of the actions that lead to it. So, the proof constructs a function which maps each value in $\kappa(A)$ to a real number in such a way that, when computing the order of a history, it suffices adding up the orders of each of the actions that lead to it. Still, there are, in general, many likelihood assignments that generate each given consistent assessment. One can then ask whether consistent assessments are generated by likelihood assignments with integer orders.

We say that a likelihood assignment (λ, f) is an *integer likelihood assignment* if $\lambda(a) \in \mathbb{Z}_+$ for all $a \in A$. The following corollary to Theorem 3.1, proven in Appendix

A, establishes the each consistent assessment is generated by an integer likelihood assignment.

Corollary B.1. *An assessment is consistent if and only if it is generated by some integer likelihood assignment.*

B.2 Lexicographic Probability Systems

We now compare likelihood assignments and LPSs by following an approach similar to Govindan and Klumpp (2003), that is, we consider LPSs over the set of strategy profiles (not over the beliefs of each player about the beliefs of the other players). To simplify the exposition, assume that each information set is associated to only one player.

In a normal-form game, an LPS is defined as a finite sequence of distributions over the space of pure strategies S. In order to avoid additional independence conditions extensive-form games, we introduce the concept if a behavior LPS as a sequence of behavior strategies.

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

Let $S \equiv \prod_{h\in H} A^h$, interpreted as the set of pure strategies of the game. We can then associate an LPS of order $\hat{L}\,|H|$ to a behavior LPS by defining, for each order $\ell = 0, ..., \hat{L}\,|H|$, the mixed strategy $\hat{\sigma}^\ell \in \Delta(S)$ as

$$\hat{\sigma}^\ell(s) \equiv C^\ell \sum_{\{(\ell_h)_h|\sum_h \ell_h=\ell\}} \prod_{h\in H} \sigma^{\ell_h}(s_h) \quad \forall s \in S \ ,$$

where C^ℓ is the appropriate constant to make σ^ℓ a probability distribution over S. We use $\hat{\lambda}(s)$ to denote the lowest order ℓ where $\hat{\sigma}^\ell(s) > 0$. The following result illustrates the relationship between likelihood assignments and (behavior) LPSs:

Theorem B.2. *If $(\sigma^\ell)_{\ell=0}^{\hat{L}}$ is a behavior LPS then $(\hat{\lambda}, \hat{f})$, with $\hat{f}(a) \equiv \sigma^{\hat{\lambda}(a)}(a)$ for all $a \in A$, is an integer likelihood assignment. Furthermore, for all $t \in T$,*

$$\hat{\lambda}(t) = \min\{\, \hat{\lambda}(s) \,|\, t \textit{ is on the path of } s \,\} \quad \textit{and}$$
$$\hat{f}(t) = \hat{\sigma}^{\hat{\lambda}(t)}(\{\, s \,|\, t \textit{ is on the path of } s \,\}) \ .$$

Proof. It is clear that if $(\sigma^\ell)_{\ell=0}^{\hat{L}}$ is a behavior LPS then $(\hat{\lambda}, \hat{f})$, with $\hat{f}(a) \equiv \sigma^{\hat{\lambda}(a)}(a)$ for all $a \in A$, is an integer likelihood assignment. Fix some history $t = (a_j)_{j=1}^J$, and let $H(t) \equiv \{h^{a_j}|j = 1, ..., J\}$. Let s be such that $\hat{\lambda}(s)$ is minimal among the strategies where

t is on path of s. Then, there exists some vector $(\ell_h)_{h\in H}$ such that $\sum_{h\in H}\hat{\lambda}(s_h)=\hat{\lambda}(s)$ and $\sigma^{\ell_h}(s_h)>0$. It is clear that, given the minimal property of s, it is the case that $\hat{\lambda}(a)=0$ for all a such that $h^a\notin H(t)$. Then, necessarily, $\hat{\lambda}(s)=\hat{\lambda}(t)$, which proves the second result of the theorem. Finally, notice that

$$\hat{\sigma}^{\hat{\lambda}(t)}(\{\,s\,|\,t \text{ is on the path of } s\,\})=\sum_{s\in S(t)}\prod_{h\in H}\sigma^{\hat{\lambda}(s_h)}(s_h)=\prod_{h\in H(t)}\hat{f}(s_h)=\hat{f}(t)\;.$$

where $S(t)$ is the set of pure strategy where t is on their path. This concludes the proof. □

Theorem B.2 is illustrative of how a likelihood assignment can be interpreted as retaining, from a given behavior LPS, the information necessary to determine the relative probabilities between histories and, as a result, it is sufficient to determine the consistency of assessments. The result implies that, for a fixed behavior LPS, the likelihood order of a history is determined by the lowest likelihood order among the strategy profiles where such a history happens with positive probability. So, histories which have a lower likelihood order are perceived as "infinitely more likely" than histories with a higher likelihood order. Furthermore, the relative weight between two histories of the same order is given by the probabilities that the mixed strategy of that order assigns to the respective sets of strategies that have each of them on path. So, a likelihood assignment selects the information that is relevant to compare likelihoods of histories so, as we see in Section 4, they are sufficient to characterize sequential equilibria.

In applications, the order $\hat{L}$ of an (behavior) LPS supporting a given assessment (μ,σ) may be large.[10] Still, as we argued, most of the components of an LPS are not necessary to determine the incentives of the players in the game. As a result, the reduced dimensionality of likelihood assignments ($\mathbb{R}^{2\,|A|}$ instead of $\mathbb{R}^{(\hat{K}+1)\,|A|}$ in a behavior LPS), their additive and multiplicative properties and the fact that they can be defined in each information set independently, make them easier to work with in applications.

Remark B.1. Even though, as Corollary B.1 establishes, all consistent assessments are generated by integer likelihood assignments, obtaining them in applications with complex dynamic games may be difficult. The reason is that, given that the orders of some histories have to match in information sets that assign them positive probability, finding

[10]It is important to notice that, in the same way that the set of orders used in an integer likelihood assignment (i.e., $\lambda(A)$) may contain sizable gaps (that is, in general, $\max(\lambda(A))$ may be much bigger than $|\lambda(A)|$), $\hat{L}$ may be (much) bigger than the maximum number of actions available in an information set in an LPS. This reason is that, in both cases, some assessments require assigning the same likelihood order to different histories composed by actions played by different players and of different likelihood orders, which generates a system of equations which, sometimes, has high-valued solutions.

integer orders for each action that led to them (so that the sum of the orders is the same) may require using large natural numbers, which may not be easy to obtain. This highlights an advantage of using likelihood assignments: the real-valued orders give flexibility to the assignment of orders, which may ease their obtention, specially when one has an economic intuition about the likelihoods with which actions are played.

B.3 Conditional Probability Systems (CPS)

A different approach to model consistent behavior is through the use of conditional probability systems, like Battigalli (1996), or relative probabilities, like Kohlberg and Reny (1997). In this section we focus on the relationship of between conditional probability systems and likelihood assignments.

We follow Battigalli (1996) in the definitions of this subsection. A *conditional probability system (CPS)* on S is 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 $S_1, S_2, S_3 \subset S$,

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

In order to ensure that a CPS is consistent with independent randomizations, one should impose the "independence property". We say that a CPS P *satisfies the independence property* if for any partition $\{H', H''\}$ of H, and two sets $S_1' \times S_1'', S_2' \times S_2'' \in (\prod_{h \in H'} A^h) \times (\prod_{h \in H''} A^h)$, we have

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

Notice that, instead of defining the independence property over the pure strategies, we define it over the space of actions (that is, we allow different information sets of the same player to be in different elements of the partition $\{H', H''\}$ of H.) In our context of an extensive-form game with a focus on behavior strategies this is without loss of generality, since each player randomizes independently in each of her information sets. Finally, a *strategic extended assessment* is a triple (σ, μ, P), where (σ, μ) is an assessment and P is a CPS satisfying

$$\mu^h(t) = P\big(S(t) \big| S(h)\big) \text{ and } \sigma(a) = P\big(\, \{ s \in S(h^a) \mid s_{h^a} = a \} \,\big|\, S(h^a)\,\big) \qquad \text{(B.1)}$$

for all $h \in H$, $t \in h$ and $a \in A$, where $S(t)$ is the set of elements in S that contain all actions that lead to t, and $S(h) = \cup_{t' \in h} S(t')$.

Now fix some likelihood assignment (λ, f). For any element $s \equiv (s_h)_{h \in H} \in S$ we define, with some abuse of notation, its likelihood order and intensity as:

$$\lambda(s) \equiv \sum_{h \in H} \lambda(s_h) \text{ and } f(s) \equiv \prod_{h \in H} f(s_h) \ .$$

Also, for a given set $S_1 \subset S$, we denote $\lambda(S_1) \equiv \min\{\lambda(s)|s \in S_1\}$ and $(S_1)_\lambda \equiv \{s \in S_1|\lambda(s)=\lambda(S_1)\}$. Finally, we say that P is the conditional probability system generated by (λ, f) if, for all $S_1, S_2 \subset S$ with $S_2 \neq \emptyset$:

$$P(S_1|S_2) \equiv \begin{cases} \frac{\sum_{s\in(S_1\cap S_2)_\lambda} f(s)}{\sum_{s\in(S_2)_\lambda} f(s)} & \text{if } \lambda(S_1 \cap S_2) = \lambda(S_2) \\ 0 & \text{otherwise.} \end{cases}$$

The following result illustrates the connection of likelihood assignments and CPS:

Theorem B.3. *Let P and (μ, σ) be, respectively, the conditional probability system and the assessment generated by some numbering (λ, f). Then, P satisfies the "independence property" and (P, μ, σ) is a strategic extended assessment.*

Proof. It is trivial to show that P satisfies the independence property. To prove that (P, μ, σ) is a strategic extended assessment, notice that for any history $t = (a_j)_{j=1}^J$,

$$\sum_{s\in S(t)_\lambda} f(s) = f(t) \ . \tag{B.2}$$

This result holds because a pure strategy s is such that $s \in S(t)_\lambda$ if and only if (1) t on its path of play (i.e., $a_j = s_{h^{a_j}}$) and, (2) if $h \notin \{h^{a_j}|j = 1, ..., J\}$, we have $s_h \in A^h_\lambda$.[11] Similarly, we have

$$\sum_{s\in S(h)_\lambda} f(s) = \sum_{t\in h_\lambda} \sum_{s\in S(t)_\lambda} f(s) = \sum_{t\in h_\lambda} f(t) \ . \tag{B.3}$$

As a result, using the fact that (μ, σ) is generated by (λ, f) and the second part of Definition 3.2, the first condition in equation (B.1) holds. To prove the second condition, fix an action $a \in A$. Using $S(a)$ to denote $\{s \in S(h^a) \mid s_{h^a} = a\}$, and an argument similar to the one used to obtain equations (B.2) and (B.3), we have

$$\sum_{s\in S(a)_\lambda} f(s) = f(a) \sum_{t\in h^a_\lambda} f(t) \ . \tag{B.4}$$

Then, the second condition in equation (B.1) holds, and the proof is done. □

Corollary 3.1 in Battigalli (1996) shows that an assessment (μ, σ) is consistent only if there is some CPS P satisfying the independence property such that (μ, σ, P) is an strategic extended assessment. Theorem B.3 provides us with one of such CPSs: it is the one generated by a likelihood assignment that generates (μ, σ).

[11] Notice that, in expression (B.2), since there is a sum over all strategies in $S(t)_\lambda$, in each information set h the sum runs over all $f(a)$ for $a \in A^h_\lambda$, which is equal to 1.

Remark B.2. Kohlberg and Reny (1997) use as a primitive object the assessment of an "outside observer" to characterize the relative probability of any $|N|$-tuple of (normal-form game) pure strategies of the players. Our construction can also be interpreted as generating, for each pair of histories of the extensive-form game, an outside observer's assessment about their relative likelihood. In fact, Theorem 3.1 shows that it is indeed the case: if $(\mu_n, \sigma_n)_n$ is a sequence of fully-mixed assessments converging to some assessment (μ, σ) if and only if there is a likelihood assignment (λ, f) that generates (μ, σ) where[12]

$$\forall t, t' \in T, \quad \lim_{n\to\infty} \frac{\Pr(t|\sigma_n)}{\Pr(t'|\sigma_n)} = \begin{cases} 0 & \text{if } \lambda(t) > \lambda(t'), \\ f(t)/f(t') & \text{if } \lambda(t) = \lambda(t'), \\ \infty & \text{if } \lambda(t) < \lambda(t'). \end{cases} \tag{B.5}$$

Still, CPSs and relative probabilities are difficult to use in applications. In our view, the main reasons are their high dimensionality ($2^{2^{|S|}}$) and the requirement that a CPS satisfies the independence property, which is sometimes difficult to verify.

[12]In the same way that different sequences supporting a given consistent assessment give different limit relative probabilities between histories, different likelihood assignments generating the same assessment may also differ on the relative probability they assign to different histories. Still, in all of them, (B.5) holds whenever t and t' belong to the same information set.