

SPECIFYING NODES AS SETS OF ACTIONS

Peter A. Streufert
Department of Economics
University of Western Ontario

ABSTRACT.

The nodes of an extensive-form game are commonly specified as sequences of actions. Rubinstein calls such nodes histories. We find that this sequential notation is superfluous in the sense that nodes can also be specified as *sets* of actions. The only cost of doing so is to rule out games with absent-minded agents.

Further, we present an application. We take an arbitrary assessment and define its infinitely-more-likely relation. This relation compares nodes, which are now specified as sets of actions. We find that if the assessment is consistent, then its infinitely-more-likely relation can be additively represented by a density function assigning numbers to actions. This construction is unexpectedly intuitive because it closely resembles that of an ordinary probability density function. Essentially, we have found a clear way to say that consistency embodies the “stochastic independence” of zero-probability actions.

Finally, we show that this intuitive result is the driving force behind the algebraic characterizations of consistency in Kreps and Wilson (1982, Appendix) and Perea, Jansen, and Peters (1997). Thus one can comfortably study consistency without reference to sequences of assessments.

Date: March 16, 2012.

I thank Audra Bowlus, Ian King, Val Lambson, Bart Lipman, Andrew McLennan, Bernhard von Stengel, Bob Wilson, and Charles Zheng; seminar participants at Auckland, Brigham Young, the London School of Economics, New South Wales, Sydney, and Western Ontario; and conference participants at the Econometric Society meeting in Sydney, the game theory conference in Stony Brook, and the Midwest mathematical economics meeting in Ann Arbor.

1. INTRODUCTION

1.1. SET-TREE GAMES

In order to define an extensive-form game, one sometimes begins with a tree consisting of nodes and edges. One then uses that tree as a skeleton on which to define actions, information sets (i.e. agents), players, chance probabilities, and payoffs. By assumption, the tree must have a distinguished node, called the initial node, which is connected to every other node by exactly one path. This node-and-edge formulation can be traced to Kuhn (1953, Section 1) and it appears today in Mas-Colell, Whinston, and Green (1995, page 227).

Node-and-edge notation is complicated, even in the clean presentation of Mas-Colell, Whinston, and Green (1995). To simplify notation, Rubinstein begins with actions rather than nodes-and-edges, and then constructs each node as the sequence of actions leading to it. Accordingly, his tree is a collection of action sequences (i.e. histories) of the form (a_1, a_2, \dots, a_N) , and his initial node is the empty sequence $\{\}$. He assumes that if (a_1, a_2, \dots, a_N) is in the tree, then $(a_1, a_2, \dots, a_{N-1})$ must also be in the tree. Hence he implicitly guarantees that the initial node is connected to every other node by exactly one path. This sequence-tree formulation appears in Osborne and Rubinstein (1994, page 200).

In this paper, we go one step further and identify each node with the *set* of actions leading to it. In particular, we define a “set tree” to be a collection of sets, which has the property that every nonempty set in the tree has a unique element whose removal results in another set of the tree. This unique element is defined to be the set’s “last action.”

It is incumbent upon us to demonstrate the sense in which such a set tree is equivalent to a sequence tree. Toward this end, we define an isomorphism between sequence trees and set trees: we say that a sequence tree is “isomorphic” to a set tree if there is an invertible map from sequences to sets, such that removing the last action of any sequence corresponds to removing the last action of the corresponding set. In this manner the isomorphism formalizes the resemblance between the concatenation of sequences and the union of sets.

Finally we define “agent recall” to mean the absence of an absent-minded agent. This condition is weaker than perfect recall. Theorem A then shows that sequence-tree games with agent recall are equivalent to set-tree games. To be precise, every sequence-tree game with agent

recall is isomorphic to exactly one set-tree game. Conversely, every set-tree game is isomorphic to exactly one sequence-tree game, and that sequence-tree game has agent recall.

Theorem A may seem implausible because a sequence specifies order and thus has more structure than a set. Some intuition can be gained from three observations. First, we assume that actions are agent-specific, without loss of generality and like many authors. As a result, the agent taking an action can be determined without knowing the order in which the actions are played. Second, we find that agent recall rules out sequences which repeat actions. Such sequences could not be faithfully represented as sets. Third, we use an example to suggest that if two actions can be played in two different orders, then there must be a previous action that determines the order.

Although the above remarks might assist with intuition, the theorem's proof remains nontrivial. First consider going from a sequence tree to a set tree. This direction is relatively easy in the sense that each sequence's set is uniquely determined as the set of actions appearing in the sequence. However, nontrivial issues remain: one must show that this map from sequences to sets is invertible, that each set has a unique last action, and that a set's last action appears as the last element of the sequence that generated the set.

Second, consider constructing a sequence tree from a set tree. This direction is relatively difficult in the sense that both uniqueness and existence become nontrivial. Uniqueness seems unlikely because a given set can be ordered as a sequence in many different ways, and, to compound matters further, the theorem admits sequences that repeat actions when it admits arbitrary sequence trees that need not satisfy agent recall. Existence is also nontrivial because sequences must be assigned to sets in such a way that the concatenation of sequences is isomorphic to the union of sets, and hence, assigning a sequence to any one set places restrictions on the assignments at all the set's subsets and supersets. Essentially, the uniqueness result shows that a set tree has a surprising amount of structure, and the existence result shows that that structure is never strong enough to prevent the construction of a sequence tree.

To our knowledge, this is the first paper to simplify games by means of set trees. We believe that this simplification will pay substantial

dividends. One such dividend appears in the second half of this paper, to which we now turn.

1.2. REPRESENTING PLAUSIBILITY BY A DENSITY FUNCTION

In Kreps and Wilson (1982), the definition of consistency incorporates some sort of stochastic independence among the agents' behavioural strategies. To see this in detail, recall that an "agent" is a synonym for an "information set", and that an "assessment" lists each agent's belief over its constituent nodes and also its behavioural strategy over its feasible actions. For any assessment, the probability of reaching any node can be calculated by first finding the set of actions leading to it, and then multiplying the probabilities that were assigned to these actions by the assessment's behavioural strategies. This calculation presumes that the behavioural strategies are stochastically independent.

Further, if every behavioural strategy in an assessment has full support, every node is reached with positive probability. Then, if every belief in such an assessment is the conditional probability distribution derived from these positive probabilities, the assessment is said to be "full-support Bayesian". A "consistent" assessment is the limit of a sequence of full-support Bayesian assessments. Since every full-support Bayesian assessment presumes that the behavioural strategies are stochastically independent, every consistent assessment must inherit something similar.

One naturally hopes to directly understand this limiting sort of stochastic independence without reference to the convergent sequence of full-support Bayesian assessments. This is subtle. If a consistent assessment specifies pure behavioural strategies, it is difficult to directly understand how these degenerate distributions are stochastically independent without reference to a convergent sequence of full-support distributions. Nonetheless much progress has been made. Blume, Brandenburger, and Dekel (1991a, Section 7), Hammond (1994, Section 6.5), and Halpern (2010, Section 6) all formulate independence in terms of non-Archimedean probability numbers. Battigalli (1996, Section 2) formulates independence in terms of conditional probability systems. Kohlberg and Reny (1997, Section 2) formulate independence in terms of relative probability systems.

We introduce a complementary perspective that is simpler. In brief, Theorem B shows that the infinite-relative-likelihood (i.e., “plausibility”) relation of a consistent assessment can be represented by a density function. To our knowledge, the literature has not analyzed infinite relative likelihood from the perspective of representation theory. This new approach works because Theorem A allows us to consider nodes as sets. Theorem B’s formulation and proof are simple in that they use nothing more than linear algebra, and they are intuitive because of surprisingly tight analogies with the foundations of ordinary probability theory.

This and the next three paragraphs provide some more details. For expository ease, we use “more plausible” as a synonym for “infinitely more likely”. Formally, we define the “plausibility relation” \succsim of an arbitrary assessment. This relation compares nodes in only two circumstances. (a) Suppose two nodes belong to the same agent (i.e., information set). Then the two are equally plausible if both are in the support of the agent’s belief, and the first is more plausible than the second if the first is in the support while the second is not. (b) Suppose one node immediately precedes another. Then the two nodes are equally plausible if the intervening action is played with positive probability, and the first node is more plausible than the second if the intervening action is played with zero probability. In this fashion, \succsim is derived from the beliefs and behavioural strategies of a given assessment.

By Theorem A, a node can be regarded as a set of actions. Analogously in probability theory, an event is a set of states. Further, a plausibility relation \succsim compares nodes. Analogously, Kraft, Pratt, and Seidenberg (1959) take as primitive a probability relation that compares events. Their probability relation embodies the notion that one event is regarded as more probable than another. They then show that a well-behaved probability relation can be represented by a density function. In detail, a density function assigns probability numbers to states. Then the probability of an event can be calculated as the sum of the probability numbers assigned to its states. Finally, the probabilities of all the events represent the original probability relation.

Analogously, Theorem B shows that a consistent assessment’s plausibility relation can be represented by a density function. In detail, a density function π assigns “plausibility” numbers to actions. Then

the “plausibility” of a node can be calculated as the sum of the plausibility numbers assigned to its actions. Finally, the plausibilities of all the nodes represent the assessment’s plausibility relation. Both this construction and the analogous one for probability are straightforward. Neither requires more than linear algebra.

Although these constructions are very similar, plausibility numbers are nonpositive rather than nonnegative. This happens because plausibility diminishes as actions accumulate. More precisely, representation and part (b) in the definition of \succsim together require that each positive-probability action be given zero plausibility and that each zero-probability action be given negative plausibility. Accordingly, a node’s plausibility (that is, the sum of the plausibility numbers of a node’s actions) is a measure of how far the node lies below the equilibrium path. It is slightly more sophisticated than (the negative of) the number of the node’s zero-probability actions because each zero-probability action can be assigned its own negative plausibility number. However, these plausibility numbers can vary only with actions and nothing else. In particular, they cannot vary with the different contexts in which a zero-probability action might be played. This invariance is a sort of stochastic independence among the zero-probability actions played by different agents.

We conclude by drawing two corollaries from Theorem B. Corollary 1 further develops our understanding of stochastic independence. Ordinarily, stochastic independence states that a joint probability density on a Cartesian product is the product of the marginal densities on the coordinate sets. The collection of nodes is not a Cartesian product. Yet, it can be embedded within a Cartesian product whose coordinates are indexed by the agents. More precisely, we map each node to the vector that lists for each agent either (a) an action of the agent that belongs to the node or (b) a “null” action which conveys the fact that the agent did not move prior to the node. Corollary 1 then shows that the embedded plausibility relation of a consistent assessment has a representation which is additive across agents. Thus, a consistent assessment’s stochastic independence across agents is very similar to a preference relation’s additive separability across consumption goods. The latter concept is due to Debreu (1960) and Gorman (1968).

Finally, while our representation approach is new, numbers resembling plausibility numbers appear in the algebraic characterizations of consistency developed by Kreps and Wilson (1982) and Perea y Mon-suwé, Jansen, and Peters (1997). These insightful characterizations are useful for calculations and as a basis for further results. Our introduction of representation theory clarifies these characterizations and substantially simplifies their proofs. Most notably, Corollary 3 fills a critical gap in the proof of Kreps and Wilson (1982, Lemma A1).

1.3. ORGANIZATION

Section 2 concerns set-tree games. Its Theorem A is proved in Appendix A.

Section 3 concerns the representation of plausibility by a density function. Its Theorem B, as well as Theorem B's two corollaries, are proved in Appendix B.

Section 4 concludes.

2. SET-TREE GAMES

This section is the first half of the paper. Here we show how each of a game's nodes can be formally identified with the *set* of actions taken to reach it.

2.1. REVIEWING SEQUENCE-TREE GAMES

We begin by reviewing Osborne and Rubinstein (1994, page 200)'s formulation of an extensive-form game. For the purposes of this paper, we call their formulation a “sequence-tree game” because it incorporates the observation that each of a game's nodes can be identified with the sequence of actions leading to it. Osborne (2008, Section 3) credits Rubinstein with this observation. We take the liberty of restating their formulation using terminology upon which we can easily build.

Let A be a set of *actions*. Then let $\bar{t} = \langle \bar{t}_n \rangle_{n=1}^{N(\bar{t})}$ denote a finite sequence of such actions, in which $N(\bar{t})$ is the length of the sequence.¹ By convention, the empty set $\{\}$ is a sequence of actions of length zero. Further, for any nonempty \bar{t} and any $0 < m \leq N(\bar{t})$, let ${}_1\bar{t}_m$ denote the sequence $\langle \bar{t}_n \rangle_{n=1}^m$. By convention, ${}_1\bar{t}_0$ equals $\{\}$ regardless of \bar{t} .

¹Osborne and Rubinstein (1994) also consider infinite sequences and hence infinite trees.

Let a *sequence tree* (A, \bar{T}) be a set A of actions together with a finite set \bar{T} of finite sequences \bar{t} of actions such that $|\bar{T}| \geq 2$, such that

$$(1) \quad (\forall \bar{t} \neq \{\}) \quad {}_1\bar{t}_{N(\bar{t})-1} \in \bar{T} ,$$

and such that every action in A appears within at least one sequence in \bar{T} (this last assumption entails no loss of generality, for if it were violated we could simply remove the superfluous actions from A). We often refer to the sequences in a sequence tree as the *nodes*² of the tree.

Given a sequence tree (A, \bar{T}) , let \bar{F} be the correspondence³ from \bar{T} into A that satisfies

$$(\forall \bar{t}) \quad \bar{F}(\bar{t}) = \{ a \mid \bar{t} \oplus (a) \in \bar{T} \} .$$

where \oplus is the concatenation operator. Since every action a in $\bar{F}(\bar{t})$ can be combined with the node \bar{t} to produce the new node $\bar{t} \oplus (a)$, the set $\bar{F}(\bar{t})$ can be understood as the set of actions that are *feasible* from \bar{t} . Then, given this feasibility correspondence \bar{F} , the set of nodes \bar{T} can be partitioned into the set of *terminal* nodes, $\bar{Z} = \{ \bar{t} \mid \bar{F}(\bar{t}) = \emptyset \}$, and the set of *nonterminal* nodes, $\bar{T} \sim \bar{Z} = \{ \bar{t} \mid \bar{F}(\bar{t}) \neq \emptyset \}$.⁴ Note that \bar{F} and \bar{Z} are derived from (A, \bar{T}) .

A game will also specify a collection $\bar{H} \subseteq \mathcal{P}(\bar{T} \sim \bar{Z})$ of *agents* (i.e., information sets) \bar{h} such that \bar{H} partitions $\bar{T} \sim \bar{Z}$ and such that

$$(2a) \quad (\forall \bar{t}^1, \bar{t}^2) \quad [(\exists \bar{h}) \{ \bar{t}^1, \bar{t}^2 \} \subseteq \bar{h}] \Rightarrow \bar{F}(\bar{t}^1) = \bar{F}(\bar{t}^2) \quad \text{and}$$

$$(2b) \quad (\forall \bar{t}^1, \bar{t}^2) \quad [(\nexists \bar{h}) \{ \bar{t}^1, \bar{t}^2 \} \subseteq \bar{h}] \Rightarrow \bar{F}(\bar{t}^1) \cap \bar{F}(\bar{t}^2) = \emptyset .$$

The first of these two implications states that the same actions are feasible from any two nodes in an agent \bar{h} . This assumption is standard and leads one to write $\bar{F}(\bar{h})$ for the set of actions feasible for agent \bar{h} .⁵ The second implication states that actions are *agent-specific* in the sense that nodes from different agents must have different actions. Agent-specific actions are also assumed by Kreps and Wilson (1982),

²Osborne and Rubinstein (1994) refer to such a sequence as a “history” and denote it by “ h ”. We reserve “ h ” for an agent (i.e., information set).

³Although this correspondence is usually denoted by “ A ”, we reserve “ A ” for the set of all actions.

⁴As a matter of convention, we denote the empty set by $\{\}$ when it is regarded as a node and denote it by \emptyset in all other contexts.

⁵As with any correspondence, the value $\bar{F}(\bar{h})$ of the correspondence \bar{F} at the set \bar{h} is defined to be $\{ a \mid (\exists \bar{t} \in \bar{h}) a \in \bar{F}(\bar{t}) \}$. This construction is particularly natural here because (2a) implies that $(\forall \bar{t} \in \bar{h}) \bar{F}(\bar{t}) = \bar{F}(\bar{h})$.

and the assumption entails no loss of generality because one can always introduce enough actions so that agents never share actions.

Further, for the purposes of this paper, let an *augmented* partition of a set S be a collection of disjoint sets whose union is S . Notice that \emptyset can belong to an augmented partition (it cannot belong to a partition).

A *sequence-tree game* $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$ is a sequence tree (A, \bar{T}) together with (a) a collection $\bar{H} \subseteq \mathcal{P}(\bar{T} \sim \bar{Z})$ of agents (i.e., information sets) \bar{h} such that \bar{H} partitions $\bar{T} \sim \bar{Z}$ and satisfies (2), (b) a collection $\bar{I} \subseteq \mathcal{P}(\bar{H})$ of *players* \bar{i} such that \bar{I} is an augmented partition of \bar{H} , (c) a *chance* player $\bar{i}^c \in \bar{I}$, (d) a function $\bar{\rho}: \bigcup_{\bar{h} \in \bar{i}^c} \bar{F}(\bar{h}) \rightarrow (0, 1]$ which assigns a positive probability to each chance action $a \in \bigcup_{\bar{h} \in \bar{i}^c} \bar{F}(\bar{h})$, and (e) a function $\bar{u}: (\bar{I} \sim \{\bar{i}^c\}) \times \bar{Z} \rightarrow \mathbb{R}$ which specifies a *payoff* $\bar{u}_{\bar{i}}(\bar{t})$ to each nonchance player $\bar{i} \in \bar{I} \sim \{\bar{i}^c\}$ at each terminal node $\bar{t} \in \bar{Z}$. By assumption, the chance probabilities are assumed to satisfy $(\forall \bar{h} \in \bar{i}^c) \sum_{a \in \bar{F}(\bar{h})} \bar{\rho}(a) = 1$ so that they specify a probability distribution at each chance agent $\bar{h} \in \bar{i}^c$.

Note that an empty player $\bar{i} = \emptyset$ has no agents and no actions. Accordingly, a game “without chance” can be specified by setting the chance player $\bar{i}^c = \emptyset$. We assume without loss of generality that every nonchance player is nonempty.

2.2. DEFINING SET-TREE GAMES

This subsection introduces a new formulation of game in which the game’s nodes are sets rather than sequences.

Given a set A of actions, let $T \subseteq \mathcal{P}(A)$ be a collection of *nodes* t . Note that each node t is a subset of A , and thus nodes have been specified as sets of actions. Further, given such an (A, T) , let a *last action* of a node t be any action $a \in t$ such that $t \sim \{a\} \in T$. In other words, a last action of a node is any action in the node whose removal results in another node.

Figures 1, 2, and 3 provide three examples. In each case, the figure’s caption fully defines (A, T) , and accordingly, the definition is complete without the illustration itself. Each illustration links two nodes with an action-labelled line exactly when (a) that action is a last action of the larger set and (b) the smaller set is the larger set without that action. For example, f is the only last action of $\{e, f\}$ in Figure 1, and both f and g are last actions of $\{f, g\}$ in Figure 2.

A *set tree* (A, T) is a set A and a collection $T \subseteq \mathcal{P}(A)$ such that $|T| \geq 2$, such that $A = \bigcup T$, and such that

- (3) every nonempty $t \in T$ has a unique last action.

For example, Figure 1 fails to define a set tree because the node $\{f, g\}$ does not have a last action, and Figure 2 fails to define a set tree because the node $\{f, g\}$ has two last actions. In contrast, Figure 3 does define a set tree. Finally, note that the assumption $A = \bigcup T$ entails no loss of




FIGURE 1. $A = \{e, f, g\}$ and $T = \{\{\}, \{e\}, \{e, f\}, \{e, g\}, \{f, g\}\}$ violate assumption (3) since $\{f, g\}$ does not have a last action.




FIGURE 2. $A = \{f, g\}$ and $T = \{\{\}, \{f\}, \{g\}, \{f, g\}\}$ violate assumption (3) since $\{f, g\}$ has two last actions.




FIGURE 3. The set tree (A, T) defined by $T = \{\{\}, \{d_1\}, \{r_1\}, \{r_1, d_2\}, \{r_1, r_2\}\}$ and $A = \bigcup T$.

generality because $A \supseteq \bigcup T$ by construction and because $A \sim \bigcup T$ can be made empty by eliminating unused actions.

Given a set tree (A, T) , let F be the correspondence from T into A that satisfies

$$(\forall t) F(t) = \{ a \mid a \notin t \text{ and } t \cup \{a\} \in T \} .$$

Since every action a in $F(t)$ can be combined with the node t to produce a new node $t \cup \{a\}$, the set $F(t)$ can be understood as the set of actions that are *feasible* from t . Then, given F , the set of nodes T can be partitioned into the set of *terminal* nodes, $Z = \{ t \mid F(t) = \emptyset \}$, and the set of *nonterminal* nodes, $T \sim Z = \{ t \mid F(t) \neq \emptyset \}$. In this fashion F and Z are derived from (A, T) .

A set-tree game will also specify a collection $H \subseteq \mathcal{P}(T \sim Z)$ of *agents* (i.e., information sets) h such that H partitions $T \sim Z$ and such that

$$(4a) \quad (\forall t_1, t_2) [(\exists h)\{t_1, t_2\} \subseteq h] \Rightarrow F(t_1) = F(t_2) \text{ and}$$

$$(4b) \quad (\forall t_1, t_2) [(\exists h)\{t_1, t_2\} \subseteq h] \Leftarrow F(t_1) \cap F(t_2) \neq \emptyset .$$

This assumption (4) for a set-tree game is interpreted just as assumption (2) for a sequence-tree game.

Finally, a *set-tree game* $(A, T, H, I, i^c, \rho, u)$ is a set tree (A, T) together with (a) a collection $H \subseteq \mathcal{P}(T \sim Z)$ of agents h such that H partitions $T \sim Z$ and satisfies (4), (b) a collection $I \subseteq \mathcal{P}(H)$ of *players* i such that I is an augmented partition of H , (c) a *chance* player $i^c \in I$, (d) a function $\rho : \bigcup_{h \in i^c} F(h) \rightarrow (0, 1]$ which assigns a positive probability to each chance action $a \in \bigcup_{h \in i^c} F(h)$, and (e) a function $u : (I \sim \{i^c\}) \times Z \rightarrow \mathbb{R}$ which specifies a *payoff* $u_i(t)$ to each nonchance player $i \in I \sim \{i^c\}$ at each terminal node $t \in Z$. The chance probabilities are assumed to satisfy $(\forall h \in i^c) \sum_{a \in F(h)} \rho(a) = 1$ so that they specify a probability distribution at each chance agent $h \in i^c$. Without loss of generality, every nonchance player is assumed to be nonempty.

2.3. DEFINING AN ISOMORPHISM

This subsection defines a natural isomorphism between sequence-tree games and set-tree games. Accordingly, the isomorphism switches between nodes as sequences and nodes as sets.

Let R denote the function which takes a sequence of actions to a set of actions according to

$$R(\bar{t}) = \{ \bar{t}_n \mid n \in \{1, 2, \dots, N(\bar{t})\} \} .$$

For example, $R((r, r, d)) = \{d, r\}$, which illustrates that neither the order of actions in the sequence nor the repetition of actions in the sequence effects the value of R . The symbol “ R ” is natural in several senses. First, the set $R(\bar{t})$ is the “ R ”ange of the sequence \bar{t} . Second, R “ R ”educes a sequence to a set. And finally, R “ R ”emoves the bar as “ $R(\bar{t}) = t$ ” suggests.

A sequence tree (A, \bar{T}) is *isomorphic* to a set tree (A, T) if

$$(5a) \quad R|_{\bar{T}} \text{ is an invertible function from } \bar{T} \text{ onto } T, \text{ and}$$

$$(5b) \quad (\forall \bar{t}^*, a, \bar{t}) \quad \bar{t}^* \oplus (a) = \bar{t} \Leftrightarrow a \notin R(\bar{t}^*) \text{ and } R(\bar{t}^*) \cup \{a\} = R(\bar{t}).$$

By way of analogy, recall that two algebraic groups are “isomorphic” if there is an invertible function between the two groups which preserves the structure of each group’s binary relation in the structure of the other group’s binary relation. Here is something similar: $R|_{\bar{T}}$ is an invertible function between \bar{T} and T which preserves the structure of \bar{T} ’s concatenation in the structure of T ’s union, and conversely, preserves the structure of T ’s union in \bar{T} ’s concatenation.

This isomorphism between trees has many consequences. For example, suppose that (A, \bar{T}) and (A, T) are isomorphic, that \bar{F} is derived from (A, \bar{T}) , and that F is derived from (A, T) . Then by Lemma A.3(a) in Appendix A, we have that $\bar{F}(\bar{t}) = F(t)$ whenever $R(\bar{t}) = t$.

Next, let R_1 denote the function which takes an arbitrary set \bar{S}_1 of sequences into the corresponding set of sets according to⁶

$$R_1(\bar{S}_1) = \{ R(\bar{t}) \mid \bar{t} \in \bar{S}_1 \}.$$

For example, $R_1(\{(d, r, r), (d, s)\}) = \{\{d, r\}, \{d, s\}\}$. In general, if (A, \bar{T}) and (A, T) are isomorphic, we have that $R_1|_{\mathcal{P}(\bar{T})}$ is an invertible function from $\mathcal{P}(\bar{T})$ onto $\mathcal{P}(T)$, that $R_1(\bar{T}) = T$, and that $R_1(\bar{Z}) = Z$ (Lemma A.3(b) in Appendix A). In the sequel, a sequence-tree agent \bar{h} will be mapped to the set-tree agent $R_1(\bar{h}) = h$.

Further, let R_2 denote the function which takes an arbitrary set \bar{S}_2 of sets of sequences into the corresponding set of sets of sets according

⁶In common parlance, if $f: X \rightarrow Y$ and $B \subseteq X$ then $f(B)$ is understood to be $\{f(x) \mid x \in B\}$. Thus common parlance endows the symbol $f(\cdot)$ with two meanings, one for when the argument is an element of X and the other for when the argument is a subset of X . Our introducing R_1 is like dropping the second meaning of $f(\cdot)$ (so that $f(B)$ becomes undefined) and then introducing the symbol $f_1(\cdot)$ (so that $f_1(B)$ becomes defined). We do not use the f_1 notation in general. For example, we write $F(h)$ rather than $F_1(h)$.

to

$$R_2(\bar{S}_2) = \{ R_1(\bar{S}_1) \mid \bar{S}_1 \in \bar{S}_2 \} .$$

For instance, $R_2(\{ \{(d, r), (d, d)\}, \{(x, x)\} \}) = \{ \{ \{d, r\}, \{d\} \}, \{ \{x\} \} \}$. In general, if (A, \bar{T}) and (A, T) are isomorphic, then $R_2|_{\mathcal{P}^2(\bar{T})}$ is an invertible function from $\mathcal{P}^2(\bar{T})$ onto $\mathcal{P}^2(T)$. In the sequel, a sequence-tree player \bar{i} will be mapped to the set-tree player $R_2(\bar{i}) = i$.

Finally, say that $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$ and $(A, T, H, I, i^c, \rho, u)$ are *isomorphic* if (A, \bar{T}) and (A, T) are isomorphic,

$$(6a) \quad \{ R_1(\bar{h}) \mid \bar{h} \in \bar{H} \} = H ,$$

$$(6b) \quad \{ R_2(\bar{i}) \mid \bar{i} \in \bar{I} \} = I ,$$

$$(6c) \quad R_2(\bar{i}^c) = i^c ,$$

$$(6d) \quad \bar{\rho} = \rho , \text{ and}$$

$$(6e) \quad (\forall \bar{i} \neq \bar{i}^c)(\forall \bar{t} \in \bar{Z}) \bar{u}_{\bar{i}}(\bar{t}) = u_{R_2(\bar{i})}(R(\bar{t})) .$$

2.4. SHOWING THE ISOMORPHISM IS ONE-TO-ONE

This subsection contains Theorem A, which is our primary result. The theorem shows that isomorphism is a one-to-one correspondence between the collection of sequence-tree games with agent recall and the collection of set-tree games.

Before we define “agent recall”, we recall the concept of absent-mindedness defined by Piccione and Rubinstein (1997). Informally, an agent is absent-minded if the agent does not know whether it has already moved. Formally, an agent is *absent-minded* if there is a sequence which enters the agent more than once. In other words, an agent \bar{h} is absent-minded if there exist \bar{t} and $0 \leq m < n \leq N(\bar{t})$ such that $\{ {}_1\bar{t}_m, {}_1\bar{t}_n \} \subseteq \bar{h}$.

A sequence tree (A, \bar{T}) with agents \bar{H} is said to have *agent recall* if it has no absent-minded agents. In other words, agent recall is the absence of absent-mindedness. Agent recall is implied by perfect recall, and perfect recall is assumed by many authors including Kreps and Wilson (1982). Specifically, they define perfect recall as the combination of their equations (2.2) and (2.3). Their equation (2.2) is equivalent to agent recall by Lemma A.4(b) in Appendix A, and their equation (2.3) might be usefully called “player recall” as opposed to “agent recall.”

Theorem A. *Every sequence-tree game with agent recall is isomorphic to exactly one set-tree game. Conversely, every set-tree game is isomorphic to exactly one sequence-tree game, and that sequence-tree game has agent recall. (Proof: Lemmas A.7 and A.8 in Appendix A.)*

Thus the theorem shows that isomorphism constitutes a one-to-one correspondence between (a) the collection of sequence-tree games with agent recall and (b) the collection of set-tree games. Or, to put it another way, the structure of a sequence-tree game with agent recall is identical to the structure of a set-tree game.

2.5. THEOREM A'S INTUITION

The theorem may seem implausible because individual sequences have more structure than individual sets, simply because sets do not specify order. To overcome this initial reaction, we make four observations which informally support the theorem.

First, considerable power comes from specifying agent-specific actions. In particular, both sequence-tree games and set-tree games assume agent-specific actions (see (2) and (4)), and thus, the actions themselves encode the agents that take them. Although this assumption is powerful, it entails no loss of generality because one can always introduce enough actions so that each agent gets its own actions.

Second, sequences which repeat actions would be problematic, but agent recall rules them out (Lemma A.5 in Appendix A). For example,




FIGURE 4. The sequence (r, r) repeats the action r (and thereby precludes isomorphism). Accordingly, the agent $\bar{h} = \{\{\}, (r)\}$ is absent-minded, in violation of agent recall.

consider Figure 4’s sequence-tree game, which is essentially Piccione and Rubinstein (1997, Example 1). Here the sequence (r, r) repeats the action r . This repetition prevents isomorphism because the distinct sequences (r, r) and (r) map to the same set $\{r, r\} = \{r\}$, and thus $R|_{\bar{T}}$ is not invertible in violation of (5a). However, this example violates agent recall because its only agent is absent-minded: the sequence $\bar{t} = (r)$ enters the agent twice, once at ${}_1\bar{t}_0 = \{\}$ and again at $\bar{t} = (r)$.

Third, assumption (1) for a sequence tree resembles assumption (3) for a set tree. Intuitively, the “last action” of a sequence is its last element, and assumption (1) guarantees that the sequence without this “last action” is also a member of the tree. Formally, the last action of a *set* is defined to be an action whose removal results in another member of the tree, and assumption (3) guarantees that each set in the tree has exactly one last action. Essentially, the “last action” of a sequence is explicitly stated by the sequence itself, while the last action of a *set* is implicitly determined by the entire tree.

Fourth, the following example suggests that a set of actions in a set tree can only be played in one order, because if that order is not exogenous, it must have been determined by some action(s) in the set itself.

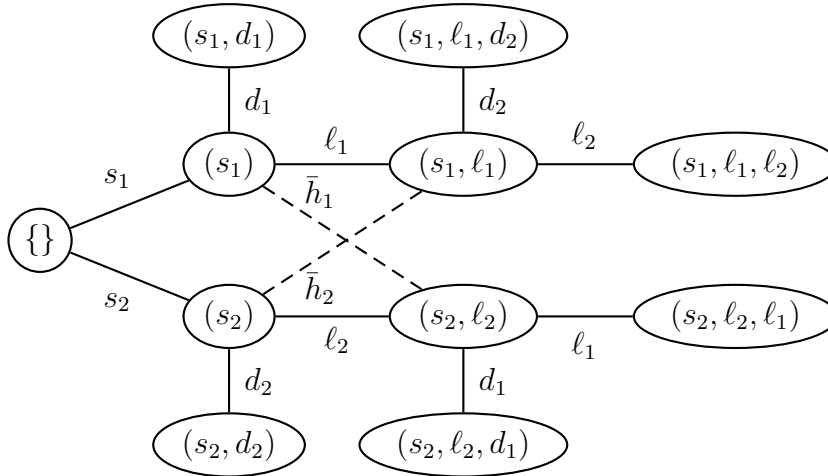


FIGURE 5. A sequence tree in which the order of actions appears to matter. The two agents $\bar{h}_1 = \{(s_1), (s_2, \ell_2)\}$ and $\bar{h}_2 = \{(s_2), (s_1, \ell_1)\}$ belong to the two enemies.

Imagine that two enemies refuse to be in the same room, and that a peace-loving mediator volunteers to carry a box from one to the other. While she imagines that the box will be used to convey a peace offering, the enemies imagine sending something offensive. For some reason, the enemies do not know who will get the box first, and thus neither enemy knows when receiving the box if the box is empty or already loaded.

Figure 5 specifies this situation using a sequence tree. The mediator decides to start with either the first enemy (s_1) or the second enemy (s_2), and then the two enemies either look (ℓ) in the box with intent to load it or decline (d) to open it. Clearly the game is all about which of the enemies moves first.

Yet, this situation can be specified as a set tree by simply turning the figure’s sequences into sets. Each set of actions can only be played in one order because any ambiguity is resolved by another action in the set. For example, the set $\{\ell_1, \ell_2, s_2\}$ can only be played in the order (s_2, ℓ_2, ℓ_1) because the set contains s_2 .

In summary, four observations make the theorem more intuitive: agent-specific actions encode agents within actions; agent recall precludes the repetition of actions; “last actions” in sequence trees resemble last actions in set trees; and order can often be encoded within sets.

2.6. THEOREM A’S NONTRIVIALITY

Now that the theorem appears more intuitive, we argue that its proof is still nontrivial. Many, but not all, of the difficulties stem from the fact that a sequence has more structure than a set. In particular, a sequence specifies the order in which moves are taken and a set does not. The example in Figure 5 only provides encouraging intuition that this gap can be bridged.

Going one direction, from sequences to sets, appears to be simple because R determines the set tree by $T = R_1(\bar{T})$ and then determines the rest of the set-tree game by (6). Additionally, Lemma A.5 (illustrated by Figure 4 above and proved in Appendix A below) simplifies matters further by showing that agent recall rules out sequences that repeat actions.

However, substantial issues of order remain. First, is $R|_{\bar{T}}$ invertible, or could the sequence tree \bar{T} have two sequences with the same actions

in different orders? Second, even if $R|_{\bar{T}}$ is invertible, could a set in T have multiple last actions, as would be the case in Figure 4, where both r and d would be last actions of $R((r, d)) = \{r, d\}$? Third, even if every set in T has a unique last action, could the last action of a set be in the middle, rather than at the end, of the sequence corresponding to the set? These issues are addressed in Lemmas A.6 and A.7 of Appendix A.

Going the other direction, from sets to sequences, is harder in the sense that one must figure out how to define the sequence tree. Both uniqueness and existence are nontrivial.

The theorem's claim about uniqueness is strong. It claims that each set tree corresponds to no more than one sequence tree, and further, that this uniqueness stands even if the candidate sequence trees are not required to satisfy agent recall. This claim is different than the claim that $R|_{\bar{T}}$ is an invertible function for any \bar{T} with agent recall. Rather, it says that for any T there is at most one \bar{T} which makes $R|_{\bar{T}}$ an invertible function onto T . This is a strong statement because the many possible ways of constructing the sequences of \bar{T} admit many possible ways of ordering the actions in the sets of T . Further, the possibility of defining a \bar{T} without agent recall admits the further possibility of defining sequences which repeat actions (Lemma A.5). Nonetheless, the implicit structure of a set tree T precludes all this. This is proved in Step 1 of Lemma A.8's proof.

Proving existence requires finding a way to assign sequences to sets in such a way that the concatenation of sequences is isomorphic to the union of sets, as specified in (5b). This is nontrivial because assigning a sequence to a set has implications for the assignments at all the set's subsets and supersets. The solution can be found in steps 2–6 of Lemma A.8's proof.

In summary, the uniqueness result shows that a set tree has a surprising amount of implicit structure. Then the existence result shows that that structure is never strong enough to prevent the construction of a sequence tree. Thus a sequence tree with agent recall explicitly spells out the implicit structure of a set tree.

3. REPRESENTING PLAUSIBILITY BY A DENSITY FUNCTION

This section is the second half of the paper. It shows that the plausibility relation of a consistent assessment can be represented by a density function.

3.1. REVIEWING KREPS-WILSON CONSISTENCY

This section is formulated in terms of a set-tree game $(A, T, H, I, i^c, \rho, u)$. By Theorem A, such a set-tree game implicitly assumes nothing more than agent recall. Agent recall is weaker than perfect recall, as discussed just before Theorem A's statement. Perfect recall, in turn, is a standard assumption in the literature on consistency ever since Kreps and Wilson (1982, pages 863 and 867).

The remainder of this subsection reformulates Kreps-Wilson consistency in terms of a set-tree game. First, we introduce notation that partitions the nodes and actions into those of the chance player and those of the strategic players. Since the set H of agents is partitioned by the set I of players, we can partition H into the set i^c of chance agents and the set $\{h|h \notin i^c\}$ of strategic (i.e. nonchance) agents. Then since the set $T \sim Z$ of nonterminal nodes is partitioned by H , we can partition $T \sim Z$ into the set T^c of chance nodes and the set T^s of strategic (i.e. decision) nodes:

$$T^c = \bigcup_{h \in i^c} h \text{ and } T^s = \bigcup_{h \notin i^c} h .$$

Similarly, since the set A of actions has the indexed partition $\langle F(h) \rangle_h$ (Lemma A.2 in Appendix A), we can partition A into the set A^c of chance actions and the set A^s of strategic actions:

$$A^c = \bigcup_{h \in i^c} F(h) \text{ and } A^s = \bigcup_{h \notin i^c} F(h) .$$

Note that T^c , T^s , A^c , and A^s are derived from the given game, and that the definition of A^c allows us to write $\rho: A^c \rightarrow (0, 1]$ rather than $\rho: \bigcup_{h \in i^c} F(h) \rightarrow (0, 1]$ as we did in Subsection 2.2.

Second, we introduce notation for strategies, beliefs, and assessments. A (behavioural) *strategy profile* is a function $\sigma: A^s \rightarrow [0, 1]$ such that $(\forall h \notin i^c) \sum_{a \in F(h)} \sigma(a) = 1$. Thus a strategy profile specifies a probability distribution $\sigma|_{F(h)}$ over the feasible set $F(h)$ of each strategic agent h . This $\sigma|_{F(h)}$ is h 's strategy. A *belief system* is a function $\beta: T^s \rightarrow [0, 1]$ such that $(\forall h \notin i^c) \sum_{t \in h} \beta(t) = 1$. Thus a belief system specifies a probability distribution $\beta|_h$ over each strategic agent h . This $\beta|_h$

is h 's belief. Finally, an *assessment* (σ, β) consists of a strategy profile σ and a belief system β .

Third, an assessment (σ, β) is *full-support Bayesian* if σ assumes only positive values and

$$(7) \quad (\forall h \in H^s)(\forall t \in h) \quad \beta(t) = \frac{\Pi_{a \in t}(\rho \cup \sigma)(a)}{\sum_{t' \in h} \Pi_{a \in t'}(\rho \cup \sigma)(a)} .$$

This equation calculates the belief $\beta|_h$ over any strategic agent h by means of the conditional probability law. Note that

$$\Pi_{a \in t}(\rho \cup \sigma)(a) = \Pi_{a \in t \cap A^c} \rho(a) \times \Pi_{a \in t \cap A^s} \sigma(a) .$$

is the probability of reaching node t . Here $\rho \cup \sigma$ is the union of the functions ρ and σ . In particular, $\rho \cup \sigma: A \rightarrow [0, 1]$ since $\rho: A^c \rightarrow [0, 1]$, since $\sigma: A^s \rightarrow [0, 1]$, and since $\{A^c, A^s\}$ partitions A . The positive values of ρ and σ imply that the denominator in (7) is always positive.

Finally, an assessment is *Kreps-Wilson consistent* if it is the limit of a sequence of full-support Bayesian assessments.

3.2. DEFINING AN ASSESSMENT'S PLAUSIBILITY RELATION \succsim

This subsection defines the plausibility relation \succsim of an arbitrary assessment (σ, β) . The assessment need not be consistent. The relation \succsim compares nodes. It is constructed from five components, in the five paragraphs that follow the next one. We introduce the word “plausibility” in lieu of the familiar phrase “infinite relative likelihood” only because it is shorter.

We illustrate this construction by repeatedly referring to the assessment (σ, β) of Figure 6. This figure's game tree is essentially that of Kreps and Ramey (1987, Figure 1). A casual interpretation of this game tree might be that you manage two workers, that each has a switch, and that a lamp turns on exactly when both switches are on. You can observe the lamp but not the switches, and then if the lamp is dark, you can choose to penalize either the first worker or the second worker. The figure also specifies an assessment (σ, β) . Casually, this might describe an equilibrium-like situation in which both workers work because (a) they think that if the light is dark, you would place probability 0.4 on only the first worker dozing, probability 0.4 on only the second worker dozing, and probability 0.2 on both workers dozing, (b) they see that this belief would induce you to randomize between

the two punishments, and (c) the threat of this randomized penalty motivates them both to work.

First, from the strategy profile σ derive the relation

$$\succsim = \{ (t, t \cup \{a\}) \mid a \in F(t) \text{ and } \sigma(a) = 0 \} .$$

As with any relation, the notations $(t^1, t^2) \in \succsim$ and $t^1 \succsim t^2$ are equivalent. Thus the definition of \succsim says that $t^1 \succsim t^2$ iff t^1 immediately precedes t^2 and the action leading from t^1 to t^2 is played by σ with zero probability. In such a case, we say that t^1 is “more plausible” than t^2 in the sense that t^1 is infinitely more likely than t^2 . For example, $\{\} \succsim \{d_1\}$, $\{d_1\} \succsim \{d_1, d_2\}$, and $\{w_1\} \succsim \{w_1, d_2\}$ in Figure 6.

Second, from the strategy profile σ derive the relation

$$\approx = \{ (t, t \cup \{a\}) \mid a \in F(t) \text{ and } \sigma(a) > 0 \}$$

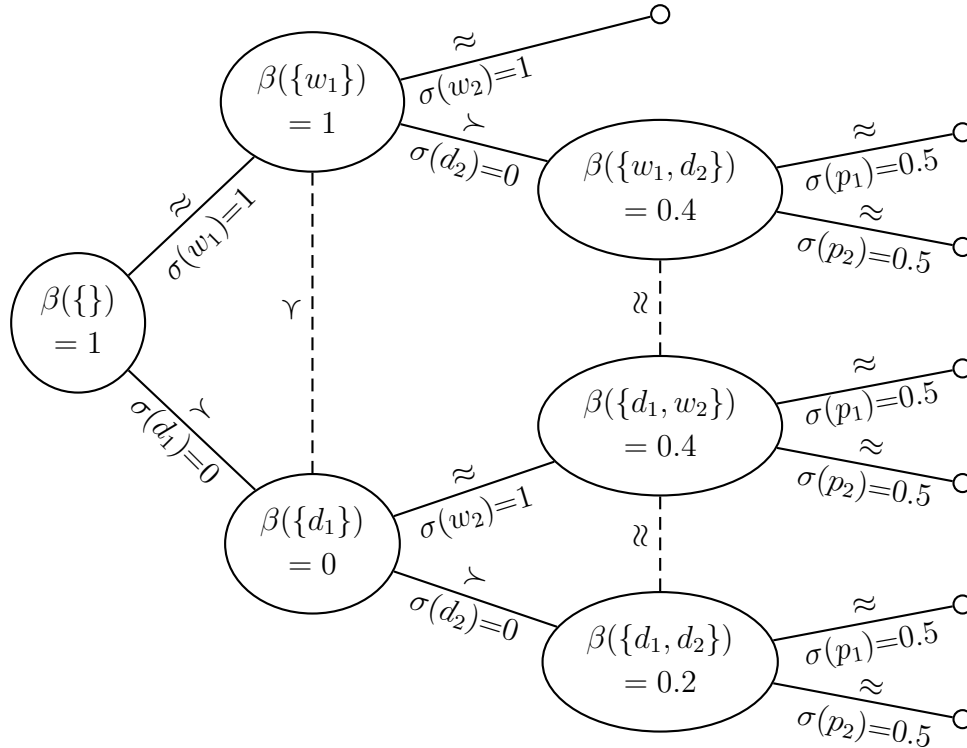


FIGURE 6. An assessment (σ, β) with its plausibility relation \succsim .

$$\cup \{ (t \cup \{a\}, t) \mid a \in F(t) \text{ and } \sigma(a) > 0 \} .$$

The definition of \approx states that both $t^1 \approx t^2$ and $t^2 \approx t^1$ hold if t^1 immediately precedes t^2 and the action leading from t^1 to t^2 is played by σ with positive probability. In such a case, we say that t^1 and t^2 are “tied in plausibility” in the sense that neither can be infinitely more likely than the other. For example, Figure 6 shows $\{ \} \approx \{w_1\}$, $\{w_1\} \approx \{w_1, w_2\}$, $\{d_1\} \approx \{d_1, w_2\}$, $\{d_1, d_2\} \approx \{d_1, d_2, p_1\}$, and five other pairs like the last one which also end in terminal nodes. (The converses of these nine pairs are also in \approx because \approx was defined to be symmetric.)

Third, this notion of tying in plausibility applies not only to strategic actions, but also to chance actions, which are played with positive probability by assumption. Accordingly, we define the relation

$$\begin{aligned} \approx^c &= \{ (t, t \cup \{a\}) \mid a \in F(t) \text{ and } a \in A^c \} \\ &\cup \{ (t \cup \{a\}, t) \mid a \in F(t) \text{ and } a \in A^c \} . \end{aligned}$$

Thus both $t^1 \approx^c t^2$ and $t^2 \approx^c t^1$ hold if t^1 is a chance node that immediately precedes t^2 . Unlike the other components of \succcurlyeq , \approx^c depends only on the game and not the assessment.

Fourth, from the belief system β derive the two relations

$$\begin{aligned} \succ^\beta &= \{ (t^1, t^2) \mid (\exists h \in H^s) \{t^1, t^2\} \subseteq h, \beta(t^1) > 0, \text{ and } \beta(t^2) = 0 \} \text{ and} \\ \approx^\beta &= \{ (t^1, t^2) \mid (\exists h \in H^s) \{t^1, t^2\} \subseteq h, t^1 \neq t^2, \beta(t^1) > 0, \text{ and } \beta(t^2) > 0 \} . \end{aligned}$$

Thus a node in the support of an agent’s belief is more plausible than any node outside the support and is tied with any other node inside the support. For example, Figure 6 shows $\{w_1\} \succ^\beta \{d_1\}$, $\{w_1, d_2\} \approx^\beta \{d_1, w_2\}$, and $\{d_1, w_2\} \approx^\beta \{d_1, d_2\}$. (The relation \approx^β also contains $(\{w_1, d_2\}, \{d_1, d_2\})$, and because the relation is symmetric, the converses of the three pairs already mentioned.)

Fifth and finally, we define \succ , \approx , and \succcurlyeq . Let \succ be the union of \succ^σ and \succ^β . Let \approx be the union of \approx^σ , \approx^c , and \approx^β . Let \succcurlyeq be the union of \succ and \approx . The following result is intuitive but not obvious.

Lemma 3.1. *Suppose that \succ^σ , \approx^σ , \approx^c , \succ^β , \approx^β , \succ , \approx , and \succcurlyeq are derived from some assessment. Then \succ is the asymmetric part of \succcurlyeq , and \succ is partitioned by $\{\succ^\sigma, \succ^\beta\}$. Similarly, \approx is the symmetric part of \succcurlyeq , and \approx is partitioned by $\{\approx^\sigma, \approx^c, \approx^\beta\}$. (Proof B.1 in Appendix B.)*

The typical plausibility relation \succsim is pervasively incomplete in the sense that it fails to compare many pairs of nodes. For instance, neither $\{\} \succsim \{d_1, d_2\}$ nor $\{d_1, d_2\} \succsim \{\}$ in Figure 6's example. Further, because of this pervasive incompleteness, the typical \succsim is also intransitive. For instance, transitivity is violated by the lack of $\{\} \succsim \{d_1, d_2\}$ in Figure 6's example.

Our \succsim differs from the infinite-relative-likelihood relations in the literature to the extent that it is derived directly from an *arbitrary* assessment. In contrast, most contributions in the literature have shown that a consistent assessment implies a rich probability structure which features infinite relative likelihoods. Such rich probability structures include the conditional probability systems of Myerson (1986), the logarithmic likelihood ratios of McLennan (1989), the lexicographic probability systems of Blume, Brandenburger, and Dekel (1991b), the non-standard probability systems of Hammond (1994) and Halpern (2010), and the relative probability systems of Kohlberg and Reny (1997). In accord with this difference, our \succsim is incomplete and intransitive while theirs are complete and transitive, and our \succsim is easier to derive because its definition bypasses their rich probability structures.

3.3. DERIVING A DENSITY-FUNCTION REPRESENTATION FOR \succsim

From an abstract perspective, \succsim is a binary relation comparing subsets $t \subseteq A$ of a space A of actions a . Similarly, Kraft, Pratt, and Seidenberg (1959) and Scott (1964) consider a binary relation \succsim comparing subsets $e \subseteq \Omega$ of a space Ω of states ω . There, the statement $e^1 \succ e^2$ means that the event e^1 is regarded as “more probable” than e^2 , and the statement $e^1 \approx e^2$ means that the events e^1 and e^2 are regarded as “equally probable”. Kraft, Pratt, and Seidenberg (1959, Theorem 2) and Scott (1964, Theorem 4.1) then state conditions on \succsim which imply the existence of a probability density function $p: \Omega \rightarrow [0, 1]$ such that for all e^1 and e^2

$$\begin{aligned} e^1 \succ e^2 &\Rightarrow \sum_{\omega \in e^1} p(\omega) > \sum_{\omega \in e^2} p(\omega) \text{ and} \\ e^1 \approx e^2 &\Rightarrow \sum_{\omega \in e^1} p(\omega) \geq \sum_{\omega \in e^2} p(\omega) . \end{aligned}$$

In brief, they find conditions under which a probability relation can be represented by a probability density function. The following theorem is similar.

Theorem B. *Consider an assessment and its plausibility relation \succsim . If the assessment is consistent, there exists $\pi:A\rightarrow\mathbb{Z}_-$ such that for all t^1 and t^2*

$$(8) \quad \begin{aligned} t^1 \succ t^2 &\Rightarrow \sum_{a\in t^1}\pi(a) > \sum_{a\in t^2}\pi(a) \text{ and} \\ t^1 \succsim t^2 &\Rightarrow \sum_{a\in t^1}\pi(a) \geq \sum_{a\in t^2}\pi(a) . \end{aligned}$$

(Proof B.5 in Appendix B. \mathbb{Z}_- is the set of nonpositive integers.)

Subsequent subsections will interpret Theorem B in the context of game theory. The remainder of this subsection sketches the theorem's proof and thereby draws further parallels with the foundations of ordinary probability theory.

To begin, consider an arbitrary finite set A and an arbitrary binary relation \succsim comparing subsets of A , which are denoted here by $s\subseteq A$ and $t\subseteq A$. Let a *cancelling sample* from \succsim be a finite indexed collection $\langle(s^m, t^m)\rangle_{m=1}^M$ of pairs (s^m, t^m) taken from \succsim such that

$$(\forall a) |\{m|a\in s^m\}| = |\{m|a\in t^m\}| .$$

Note that the sample is taken “with replacement” in the sense that a pair can appear more than once. Further, by the equation every action appearing on the left side of some pair is “cancelled” by the identical action appearing on the right side of that or some other pair. For example, if $\{a, a'\} \succsim \{a, a'\}$, then a cancelling sample from \succsim is given by $M=1$ and $(s^1, t^1) = (\{a, a'\}, \{a, a'\})$. The relation \succsim is said to satisfy the *cancellation law* if every cancelling sample from \succsim must be taken from the symmetric part of \succsim .

The cancellation law is implied by the existence of a density-function representation $\varphi:A\rightarrow\mathbb{R}$. This is easily seen if one assumes the existence of φ and considers a cancelling sample $\langle(s^m, t^m)\rangle_{m=1}^M$. By the sample's cancelling,

$$(9) \quad \sum_{m=1}^M \sum_{a\in s^m} \varphi(a) = \sum_{m=1}^M \sum_{a\in t^m} \varphi(a) .$$

By representation, $\sum_{a\in s^m} \varphi(a) > \sum_{a\in t^m} \varphi(a)$ for every (s^m, t^m) from the asymmetric part of \succsim , and $\sum_{a\in s^m} \varphi(a) = \sum_{a\in t^m} \varphi(a)$ for every (s^m, t^m) from the symmetric part of \succsim . Thus $\sum_{m=1}^M \sum_{a\in s^m} \varphi(a) = \sum_{m=1}^M \sum_{a\in t^m} \varphi(a)$ iff every pair in the sample is taken from the symmetric part of \succsim . Thus by (9) the sample must have been taken from the symmetric part of \succsim .

Interestingly, the converse also holds, and hence the cancellation law is equivalent to the existence of a density-function representation. This

result follows from Farkas' Lemma. It undergirds the intuitive foundations for probability in Kraft, Pratt, and Seidenberg (1959, Theorem 2) and Scott (1964, Theorem 4.1). It also undergirds the abstract representation theory in Krantz, Luce, Suppes, and Tversky (1971, Sections 2.3 and 9.2) and Narens (1985, pages 263-265).⁷ Lemma B.3 in Appendix B is a minor adaptation of this well-known result.

Now return to Theorem B. The plausibility relation \succsim of a consistent assessment (β, σ) must obey the cancellation law. To see this, let $\langle (\beta_n, \sigma_n) \rangle_{n=1}^\infty$ be a sequence of Bayesian full-support assessments that converge to (β, σ) , and take a cancelling sample $\langle (s^m, t^m) \rangle_{m=1}^M$ from \succsim . By the sample's cancelling,

$$(10) \quad (\forall n) \prod_{m=1}^M \frac{\prod_{a \in t^m} \rho \cup \sigma_n(a)}{\prod_{a \in s^m} \rho \cup \sigma_n(a)} = 1 .$$

Further, by the straightforward argument of Lemma B.4 in Appendix B, consistency implies (a) that

$$\lim_n \frac{\prod_{a \in t^m} \rho \cup \sigma_n(a)}{\prod_{a \in s^m} \rho \cup \sigma_n(a)} = 0$$

for every (s^m, t^m) in the asymmetric part of \succsim (note that s^m is in the denominator), and also (b) that

$$\lim_n \frac{\prod_{a \in t^m} \rho \cup \sigma_n(a)}{\prod_{a \in s^m} \rho \cup \sigma_n(a)} \in (0, \infty)$$

for every (s^m, t^m) in the symmetric part of \succsim . Thus

$$\lim_n \prod_{m=1}^M \frac{\prod_{a \in t^m} \rho \cup \sigma_n(a)}{\prod_{a \in s^m} \rho \cup \sigma_n(a)} \in (0, \infty)$$

iff the sample was taken from the symmetric part of \succsim . Thus by (10), the sample must have been taken from the symmetric part of \succsim .

Since the plausibility relation of a consistent assessment must obey the cancellation law by the last paragraph, and since the cancellation law is equivalent to the existence of a density-function representation by the paragraph before, the plausibility relation of a consistent assessment must have a density-function representation. This is the gist of Theorem B's proof.

⁷These classic results over discrete spaces complement Debreu (1960)'s derivation of an additive representation over continuum product spaces. Debreu imposes weaker cancellation assumptions (e.g., Debreu (1960, Assumption 1.3)) and compensates with topological assumptions.

In addition, the proof shows that the density function can be made to take nonpositive and integer values. The existence of an integer-valued density function is proved by using a version of Farkas' Lemma for rational matrices (Fact B.2 in Appendix B). Then, this (or any other) density function must assume nonpositive values since (a) the representation of $\approx \subseteq \succcurlyeq$ requires that zero plausibility be assigned to every action played with positive probability, and (b) the representation of $\succ \subseteq \succcurlyeq$ requires that negative plausibility be assigned to every action played with zero probability.

3.4. STOCHASTIC INDEPENDENCE

As noted in the previous paragraph, a positive-probability action is assigned a zero plausibility number and a zero-probability action is assigned a negative plausibility number. Thus a node's plausibility is a measure of how far the node is below the equilibrium path. This measure is slightly more sophisticated than (the negative of) the number of the node's zero-probability actions because each zero-probability action can be assigned its own negative plausibility number. However, a zero-probability action's plausibility number cannot vary with the context in which the action is played. That invariance is a sort of stochastic independence among zero-probability actions.

For example, consider the assessment of Figure 6. If the assessment were consistent, Theorem B would imply the existence of a plausibility density function π such that, among other things, (a) $\pi(d_2) < 0$ since the second worker dozes with zero probability, (b) $\pi(w_2) = 0$ since the second worker works with positive probability, and (c) $\pi(d_1) + \pi(d_2) = \pi(d_1) + \pi(w_2)$ since both $\{d_1, d_2\}$ and $\{d_1, w_2\}$ are in the support of the manager's belief. Since (a) and (b) together imply $\pi(d_1) + \pi(d_2) < \pi(d_1) + \pi(w_2)$, this cannot be done, and thus, the assessment is inconsistent.⁸ Note that we have implicitly assigned the same plausibility number to d_1 en route to $\{d_1, d_2\}$ as to d_1 en route to $\{d_1, w_2\}$. This invariance is a sort of stochastic independence: the plausibility number of an action cannot vary with the context in which the action is played.

⁸There is also a faster way to show the inconsistency of this assessment. A corollary of Theorem B is that a consistent assessment's plausibility relation must have an ordered extension. To prove this, one merely needs to note that one such ordered extension is represented by $\varphi(t) = \sum_{a \in t} \pi(a)$. The assessment of Figure 6 cannot have an ordered extension because $\{d_1\} \succ \{d_1, d_2\}$ and yet $\{d_1\} \approx \{d_1, w_2\} \approx \{d_1, d_2\}$.

The remainder of this subsection provides a second perspective on stochastic independence. It uses the concept of additive separability⁹ from preference theory to show that consistency implies a sort of stochastic independence across agents.

Since preferences are defined over vectors of consumption goods, the first task is to learn how to regard a node as a *vector* of actions rather than a set of actions. Accordingly, we embed the set T of nodes within a Cartesian product whose coordinates are indexed by the agents. First, create a “null” action o , and then, for each agent h , let $\dot{F}(h) = \{o\} \cup F(h)$. Thus $\dot{F}(h)$ is agent h ’s “expanded action set.” Next, consider the Cartesian product $\Pi_h \dot{F}(h)$ and let $\dot{i} = \langle \dot{i}_h \rangle_h$ denote an arbitrary vector in this product. Finally, let V be the function that maps each node $t \in T$ to the vector $V(t) \in \Pi_h \dot{F}(h)$ that is defined at each h by

$$(11) \quad [V(t)]_h = \begin{pmatrix} o & \text{if } |t \cap F(h)|=0 \\ \text{the element of } t \cap F(h) & \text{if } |t \cap F(h)|=1 \end{pmatrix} .$$

The following lemma proves that V is well-defined by showing that $|t \cap F(h)|$ is 0 or 1 for any t and h . It also proves a few other basic facts about V .

Lemma 3.2. *V is a well-defined and invertible function from T onto $V(T) \subseteq \Pi_h \dot{F}(h)$. Further, $V^{-1}(\dot{i}) = \{\dot{i}_h | \dot{i}_h \in A\}$ for every $\dot{i} \in V(T)$. (Proof B.7 in Appendix B.)*

For example, consider the set tree that is isomorphic to the sequence tree of Figure 5’s example. The three agents are

$$\begin{aligned} h^0 &= \{\} \text{ with } \dot{F}(h^0) = \{o, s_1, s_2\} , \\ h^1 &= \{\{s_1\}, \{s_2, o_2\}\} \text{ with } \dot{F}(h^1) = \{o, \ell_1, d_1\} , \text{ and} \\ h^2 &= \{\{s_2\}, \{s_1, o_1\}\} \text{ with } \dot{F}(h^2) = \{o, \ell_2, d_2\} . \end{aligned}$$

Thus the product $\Pi_h \dot{F}(h)$ has $3^3=27$ vectors. Meanwhile, the set $V(T) \subseteq \Pi_h \dot{F}(h)$ has those 11 vectors which correspond to the 11 nodes in T . Three of these vectors are

$$\begin{aligned} V(\{\}) &= (o, o, o) , \\ V(\{s_1, \ell_1\}) &= (s_1, \ell_1, o) , \text{ and} \end{aligned}$$

⁹For the additive separability of a preference relation, see Debreu (1960), Gorman (1968), and Blackorby, Primont, and Russell (1978, Section 4.4).

$$V(\{s_2, \ell_2\}) = (s_2, o, \ell_2) .$$

Note that an action's position in the vector is determined by its agent and not by the order of play.

Now consider an assessment and its plausibility relation \succsim . Let the *embedding* of \succsim in $\Pi_h \dot{F}(h)$ be the binary relation

$$\dot{\succsim} = \{ (V(t^1), V(t^2)) \mid (t^1, t^2) \in \succsim \} .$$

Since V is invertible by Lemma 3.2, there is a one-to-one correspondence between the pairs of \succsim and the pairs of $\dot{\succsim}$. For example, for any plausibility relation \succsim over the set tree of Figure 5,

$$\{s_1, \ell_1\} \succsim \{s_2, \ell_2\} \Leftrightarrow (s_1, \ell_1, o) \dot{\succsim} (s_2, o, \ell_2) .$$

Finally, we review three standard definitions for an arbitrary binary relation \succsim^* over the Cartesian product $\Pi_h \dot{F}(h)$. First, \succsim^* is an *ordering* if it is complete and transitive. Second, \succsim^* *extends* $\dot{\succsim}$ if for all vectors t^1 and t^2

$$\begin{aligned} t^1 \dot{\succsim} t^2 &\Rightarrow t^1 \succsim^* t^2 \quad \text{and} \\ t^1 \dot{\succ} t^2 &\Rightarrow t^1 \succ^* t^2 , \end{aligned}$$

where $\dot{\succ}$ and \succ^* are the asymmetric parts of $\dot{\succsim}$ and \succsim^* , respectively. Third, \succsim^* is *additively separable* if there exists $\langle \varphi_h : \dot{F}(h) \rightarrow \mathbb{R} \rangle_h$ such that for all vectors t^1 and t^2

$$t^1 \succsim^* t^2 \Leftrightarrow \Sigma_h \varphi_h(t_h^1) \geq \Sigma_h \varphi_h(t_h^2) .$$

Corollary 1. *Let $\dot{\succsim}$ be an assessment's plausibility relation, embedded in $\Pi_h \dot{F}(h)$. If the assessment is consistent, then $\dot{\succsim}$ can be extended to an additively separable ordering over $\Pi_h \dot{F}(h)$. (Proof B.8 in Appendix B).*

In other words, an assessment is consistent only if its embedded plausibility relation can be extended to an additively separable ordering. In this sense, the additivity of a consistent assessment's plausibility relation closely resembles the additive separability of a preference relation. Consistency implies additivity across agents much as separability requires additivity across consumption goods.

Since additive separability expresses preference “independence” across consumption goods, the above suggests that consistency requires an extended sort of stochastic “independence” across agents. This is intuitive. By definition, a consistent assessment is the limit of a sequence

of full-support Bayesian assessments, each of which incorporates the ordinary sort of stochastic independence across agents. Accordingly, Corollary 1 shows that the limiting notion of stochastic independence mimics additive separability.

3.5. PLAUSIBILITY NUMBERS ELSEWHERE

Although our perspective of density-function representation is new, numbers like plausibility numbers are familiar. Many papers use numbers like plausibility numbers to show consistency. Some arbitrarily chosen examples are the “error likelihoods” in Anderlini, Gerardi, and Lagunoff (2008, page 359), the “orders of probability” in Kobayashi (2007, page 525), and counting “steps off the equilibrium path” in Fudenberg and Levine (2006, Definition 3.2).

In contrast, there are few papers which go in the opposite and considerably more difficult direction. Like Theorem B, they derive numbers like plausibility numbers from consistency. The result closest to Theorem B is Kreps and Wilson (1982, Lemma A1), which assigns integers to actions in a fashion which “labels” the basis of a consistent assessment. The latter portion of this subsection identifies and repairs a critical lacuna in the proof of this valuable lemma.

Perea y Monswé, Jansen, and Peters (1997, Theorem 3.1) use a separating hyperplane to derive “order of likelihood” numbers, which then become part of their algebraic characterization of consistency. The logarithms of their numbers can be regarded as (non-integer) plausibility numbers, and accordingly, the perspective of density-function representation can be used to simplify their argument and to deepen our understanding of their characterization. Their insightful characterization is very useful because it enables one to search for consistent assessments by means of linear-programming techniques.

Hereafter let KW refer to the path-breaking work of Kreps and Wilson (1982). The remainder of this subsection identifies and repairs a gap in the proof of KW Lemma A1. Essentially, the labelling claimed by their proof can be derived as (the negative of) a plausibility density function.

First we explain how KW Lemma A1 supports the other results in their paper. KW contains three fundamental theorems in addition to its well-known definition of sequential equilibrium. Two of the theorems

concern the set of sequential-equilibrium outcomes: Theorem 2 shows that this set is generically finite, and Theorem 3 shows that it generically coincides with the set of perfect-equilibrium outcomes. These two theorems are derived from KW Theorem 1, which generically partitions the set of sequential-equilibrium assessments into a finite collection of tractable subsets. That theorem is based on KW Lemma 2, which partitions the set of consistent assessments into a finite collection of manifolds, and finally, that lemma is based on KW Lemma A1.

Next we recall definitions from KW in order to state their Lemma A1 precisely.¹⁰ In accord with KW pages 872 and 880, (a) let the *basis* b of an assessment (σ, β) be the union of σ 's support and β 's support, (b) let Ψ be the set of consistent assessments, (c) let Ψ_b be the set of consistent assessment with basis b , and (d) say that a basis is *consistent* if Ψ_b is nonempty. Then in accord with KW page 887, say that a basis is *labelled* by a function $K:A \rightarrow \mathbb{Z}_+$ if

$$(12a) \quad (\forall h)(\exists a \in F(h)) K(a) = 0 ,$$

$$(12b) \quad (\forall a) a \in b \text{ iff } K(a) = 0 , \text{ and}$$

$$(12c) \quad (\forall h)(\forall t \in h) t \in b \text{ iff } t \in \operatorname{argmin}\{J_K(t') | t' \in h\} ,$$

where $J_K:T \rightarrow \mathbb{Z}_+$ is defined (using our formulation of nodes as sets) by

$$(13) \quad J_K(t) = \sum_{a \in t} K(a) .$$

KW Lemma A1 correctly observes that a basis is consistent iff it can be labelled.

However, Streufert (2006) shows that the proof of KW Lemma A1 is flawed. In particular, the final paragraph on KW page 887 seeks to establish that any consistent basis can be labelled. It takes an arbitrary consistent basis, derives a binary relation \prec over a set of node-like objects, derives a function J which represents \prec , and then derives a function K over the set of actions. Then, the last line on page 887 claims but does not demonstrate that $J = J_K$. This equation is critical, for it is tantamount to claiming that \prec has an additive representation. Yet Streufert (2006, Subsection 3.2) shows by counterexample that this equation does not follow from their construction.

¹⁰This paragraph's concluding sentence is somewhat more general than KW Lemma A1 because their framework assumes that the chance player moves only at the outset. In their framework, our A^c would be their W , our T^c would be the set $\{\{w\} | w \in W\}$ (which consists of singleton nodes), our A^s would be their A , our $t \in T^s$ would be their $x \in X$, and our (σ, β) would be their (π, μ) .

The KW ordering \prec resembles our plausibility relation \succsim . Accordingly, our Theorem B's derivation of a density-function representation for \succsim can be used to prove their Lemma A1, as the following corollary makes explicit. Essentially, the negative of a plausibility density function π can serve as a labelling K .

Corollary 2 (= KW Lemma A1). *A basis is consistent iff it can be labelled. (Proof B.9 in Appendix B.)*

KW Theorems 2 and 3 have since been superseded by Govindan and Wilson (2001, Theorem 2.2) and Blume and Zame (1994, Theorem 4). Both these papers use abstract theorems about semi-algebraic sets.

However, for historical reasons, there is some merit in setting the record straight. Further, KW Theorem 1 continues to provide an explicit partition of the set of sequential-equilibrium assessments, and KW Lemma 2 continues to provide an explicit partition of the set Ψ of consistent assessments into the manifolds Ψ_b of $\{ \Psi_b \mid b \text{ is consistent} \}$. Although these partitions may be regarded as “long complicated construction[s] from the appendix in Kreps and Wilson (1982),” as remarked by Govindan and Wilson (2001, page 765), the partitions are less complicated when a labelling is understood as (the negative of) a density-function representation. Accordingly, arguments using these relatively explicit partitions may yet complement arguments using relatively abstract theorems about semi-algebraic sets.

4. CONCLUSION

In Section 2, we introduced the concept of a set tree, in which every node is formally identical to the set of actions which lead to it. We then proved that there is a natural one-to-one isomorphism between the collection of sequence-tree games with agent recall and the collection of set-tree games. Since agent recall is very weak, this new isomorphism allows one to simplify almost every game to its set-tree equivalent.

In Section 3, we applied this isomorphism. To set the stage, we defined the plausibility relation of an arbitrary assessment. Because of Section 2's isomorphism, we could analyze this binary relation between nodes as a binary relation between sets. In this fashion we discovered that the plausibility relation of a consistent assessment must have a density-function representation. This result was surprisingly straightforward and intuitive because of close parallels with the foundations

of ordinary probability theory. Finally, we used the result (a) to show that consistency's extended notion of stochastic independence parallels preference theory's familiar concept of additive separability, and (b) to repair a critical and yet relevant proof within Kreps and Wilson (1982).

APPENDIX A. SET-TREE GAMES AND THEOREM A

A.1. PARTITIONING ACTIONS BY AGENTS

The following parallel lemmas are unsurprising but necessary.

Lemma A.1. $\langle \bar{F}(\bar{h}) \rangle_{\bar{h} \in \bar{H}}$ is an indexed partition of A . In other words, $\{\bar{F}(\bar{h})|\bar{h}\}$ partitions A and $\bar{h} \mapsto \bar{F}(\bar{h})$ is invertible.

Proof. We begin with three observations.

(a) Each $\bar{F}(\bar{h})$ is nonempty. To see this, note \bar{H} partitions $\bar{T} \sim \bar{Z}$ by assumption, and thus each \bar{h} is a nonempty set of nonterminal nodes.

(b) If $\bar{h}^1 \neq \bar{h}^2$ then $\bar{F}(\bar{h}^1) \cap \bar{F}(\bar{h}^2) = \emptyset$. To see this, take any $\bar{h}^1 \neq \bar{h}^2$, any $\bar{t}^1 \in \bar{h}^1$, and any $\bar{t}^2 \in \bar{h}^2$. Since \bar{H} is a partition, we have $(\nexists \bar{h})\{\bar{t}^1, \bar{t}^2\} \subseteq \bar{h}$, and hence $\bar{F}(\bar{t}^1) \cap \bar{F}(\bar{t}^2) = \emptyset$ by the contrapositive of (2b). This implies $\bar{F}(\bar{h}^1) \cap \bar{F}(\bar{h}^2) = \emptyset$ because $\bar{F}(\bar{t}^1) = \bar{F}(\bar{h}^1)$ by $\bar{t}^1 \in \bar{h}^1$ and (2a), and because $\bar{F}(\bar{t}^2) = \bar{F}(\bar{h}^2)$ by $\bar{t}^2 \in \bar{h}^2$ and (2a).

(c) $\bigcup \{\bar{F}(\bar{h})|\bar{h}\} = A$. $\bigcup \{\bar{F}(\bar{h})|\bar{h}\} \subseteq A$ follows from the definition of \bar{F} . To see the converse, take any a . By assumption there exists some \bar{t} and some $m \leq N(\bar{t})$ such that $\bar{t}_m = a$. By assumption 1 applied $N(\bar{t}) - (m-1)$ times, both ${}_{1}\bar{t}_{m-1}$ and ${}_{1}\bar{t}_m$ are elements of \bar{T} . Thus since ${}_{1}\bar{t}_{m-1} \oplus (a) = {}_{1}\bar{t}_m$, we have $a \in \bar{F}({}_{1}\bar{t}_{m-1})$. Therefore since \bar{H} partitions $\bar{T} \sim \bar{Z}$, we have some \bar{h} such that ${}_{1}\bar{t}_{m-1} \in \bar{h}$ and hence $a \in \bar{F}(\bar{h})$.

$\{\bar{F}(\bar{h})|\bar{h}\}$ partitions A by observations (a)–(c). If $\bar{h} \mapsto \bar{F}(\bar{h})$ were not invertible, there would be $\bar{h}^1 \neq \bar{h}^2$ such that $\bar{F}(\bar{h}^1) = \bar{F}(\bar{h}^2)$. Since both $\bar{F}(\bar{h}^1)$ and $\bar{F}(\bar{h}^2)$ are both nonempty by observation (a), we would then have $\bar{h}^1 \neq \bar{h}^2$ such that $\bar{F}(\bar{h}^1) \cap \bar{F}(\bar{h}^2) \neq \emptyset$. This would contradict observation (b). \square

Lemma A.2. $\{F(h)\}_{h \in H}$ is an indexed partition of A . In other words, $\{F(h)|h\}$ partitions A and $h \mapsto F(h)$ is invertible.

Proof. We begin with three observations.

(a) Each $F(h)$ is nonempty. This holds because each h is a subset of nonterminal nodes.

(b) If $h_1 \neq h_2$ then $F(h_1) \cap F(h_2) = \emptyset$. To see this, take any $h_1 \neq h_2$, any $t_1 \in h_1$, and any $t_2 \in h_2$. Since H is a partition, we have $(\nexists h)\{t_1, t_2\} \subseteq h$,

and hence $F(t_1) \cap F(t_2) = \emptyset$ by the contrapositive of (4b). This implies $F(h_1) \cap F(h_2) = \emptyset$ because $F(t_1) = F(h_1)$ by $t_1 \in h_1$ and (4a), and because $F(t_2) = F(h_2)$ by $t_2 \in h_2$ and (4a).

(c) $\bigcup \{F(h)|h\} = A$. $\bigcup \{F(h)|h\} \subseteq A$ follows from the definition of F . To see the converse, take any a . By the assumption $A = \bigcup T$, there exists at least one t owning a and we may let t^* be the smallest such set. Note that a must be the last action of t^* , for if $a^* \neq a$ were its last action, $t^* \sim \{a^*\}$ would be a smaller set that also owns a . Hence $a \in F(t^* \sim \{a\})$. Further, since $t^* \sim \{a\}$ is nonterminal and H partitions the collection of nonterminal nodes, there is some h owning $t^* \sim \{a\}$. Thus by the last two sentences, $a \in F(h)$.

$\{F(h)|h\}$ partitions A by observations (a)-(c). If $h \mapsto F(h)$ were not invertible, there would be $h_1 \neq h_2$ such that $F(h_1) = F(h_2)$. Since both $F(h_1)$ and $F(h_2)$ are both nonempty by observation (a), we would then have $h_1 \neq h_2$ such that $F(h_1) \cap F(h_2) \neq \emptyset$. This would contradict observation (b). \square

A.2. SOME CONSEQUENCES OF ISOMORPHISM BETWEEN TREES

Each of this lemma's observations is used at least twice.

Lemma A.3. *The following hold whenever (A, T) is isomorphic to (A, \bar{T}) , F and Z are derived from (A, T) , and \bar{F} and \bar{Z} are derived from (A, \bar{T}) .*

(a) *If $t = R(\bar{t})$, then $F(t) = \bar{F}(\bar{t})$.*

(b) *$Z = R_1(\bar{Z})$.*

(c) *If $h = R_1(\bar{h})$, then $F(h) = \bar{F}(\bar{h})$.*

(d) *If $H = \{R_1(\bar{h})|\bar{h} \in \bar{H}\}$, then H partitions $T \sim Z$ iff \bar{H} partitions $\bar{T} \sim \bar{Z}$.*

(e) *If $I = \{R_2(\bar{i})|\bar{i} \in \bar{I}\}$ and $i^c = R_2(\bar{i}^c)$, then (6e) is equivalent to $(\forall i \neq i^c)(\forall t \in Z) u_i(t) = \bar{u}_{(R_2|_{\mathcal{P}^2(\bar{T})})^{-1}(i)}((R|\bar{T})^{-1}(t))$.*

Proof. (a) Suppose $t = R(\bar{t})$. Then by the assumed equality, by the definition of F , by manipulation, by the invertibility of $R|_{\bar{T}}$ (5a), by isomorphism condition (5b), by manipulation, and by the definition of \bar{F} ,

$$\begin{aligned} (\forall a) \quad & (t, a) \in F \\ \Leftrightarrow & (R(\bar{t}), a) \in F \\ \Leftrightarrow & a \notin R(\bar{t}) \text{ and } R(\bar{t}) \cup \{a\} \in T \end{aligned}$$

$$\begin{aligned}
&\Leftrightarrow (\exists t') a \notin R(\bar{t}) \text{ and } R(\bar{t}) \cup \{a\} = t' \\
&\Leftrightarrow (\exists \bar{t}') a \notin R(\bar{t}) \text{ and } R(\bar{t}) \cup \{a\} = R(\bar{t}') \\
&\Leftrightarrow (\exists \bar{t}') \bar{t} \oplus (a) = \bar{t}' \\
&\Leftrightarrow \bar{t} \oplus (a) \in \bar{T} \\
&\Leftrightarrow (\bar{t}, a) \in \bar{F}
\end{aligned}$$

This is equivalent to $(\forall a) a \in F(t) \Leftrightarrow a \in \bar{F}(\bar{t})$, which is in turn equivalent to $F(t) = \bar{F}(\bar{t})$.

(b) By the definition of R_1 , the definition of \bar{Z} , part (a), the invertibility of $R|_{\bar{T}}$ (5a), and the definition of Z ,

$$\begin{aligned}
R_1(\bar{Z}) &= \{ R(\bar{t}) \mid \bar{t} \in \bar{Z} \} \\
&= \{ R(\bar{t}) \mid \bar{F}(\bar{t}) = \emptyset \} = \{ R(\bar{t}) \mid F(R(\bar{t})) = \emptyset \} \\
&= \{ t \mid F(t) = \emptyset \} = Z .
\end{aligned}$$

(c) Suppose $h = R_1(\bar{h})$. Then

$$\begin{aligned}
&F(h) \\
&= F(R_1(\bar{h})) \\
&= F(\{R(\bar{t}) \mid \bar{t} \in \bar{h}\}) \\
&= \{F(R(\bar{t})) \mid \bar{t} \in \bar{h}\} \\
&= \{\bar{F}(\bar{t}) \mid \bar{t} \in \bar{h}\} \\
&= \bar{F}(\bar{h}) ,
\end{aligned}$$

where the second equality is the definition of $R_1(\bar{h})$ and the fourth follows from part (a).

(d) Assume $H = \{R_1(\bar{h}) \mid \bar{h} \in \bar{H}\}$. Then

$$\begin{aligned}
&H \text{ partitions } T \sim Z \\
&\Leftrightarrow \{R_1(\bar{h}) \mid \bar{h} \in \bar{H}\} \text{ partitions } T \sim Z \\
&\Leftrightarrow \{R_1(\bar{h}) \mid \bar{h} \in \bar{H}\} \text{ partitions } R_1(\bar{T}) \sim R_1(\bar{Z}) \\
&\Leftrightarrow \{\bar{h} \mid \bar{h} \in \bar{H}\} \text{ partitions } \bar{T} \sim \bar{Z} \\
&\Leftrightarrow \bar{H} \text{ partitions } \bar{T} \sim \bar{Z} ,
\end{aligned}$$

where the first equivalence follows from this part (c)'s assumption, the second from the invertibility of $R|_{\mathcal{P}(\bar{T})}$ by the invertibility of $R|_{\bar{T}}$ (5a) and from part (b), and the third from the invertibility of $R|_{\bar{T}}$ (5a).

(e) Assume $I = \{R_2(\bar{i}) \mid \bar{i} \in \bar{H}\}$ and $i^c = R_2(\bar{i}^c)$. Then

$$\begin{aligned} & (\forall i \neq i^c)(\forall t \in Z) u_i(t) = \bar{u}_{(R_2|_{\mathcal{P}^2(\bar{T})})^{-1}(i)}((R|_{\bar{T}})^{-1}(t)) \\ \Leftrightarrow & (\forall \bar{i} \neq \bar{i}^c)(\forall t \in Z) u_{R_2(\bar{i})}(t) = \bar{u}_{\bar{i}}((R|_{\bar{T}})^{-1}(t)) \\ \Leftrightarrow & (\forall \bar{i} \neq \bar{i}^c)(\forall \bar{t} \in \bar{Z}) u_{R_2(\bar{i})}(R(\bar{t})) = \bar{u}_{\bar{i}}(\bar{t}) \\ \Leftrightarrow & (6e), \end{aligned}$$

where the first and fourth equivalences are definitional, the second holds because of part (d)'s assumption and because $R_2|_{\mathcal{P}^2(\bar{T})}$ is invertible since $R|_{\bar{T}}$ is invertible, and the third holds because $Z = R_1(\bar{Z})$ by part (b) and because $R|_{\bar{T}}$ is invertible. \square

A.3. AGENT RECALL

Lemma A.4. *Each of the following is equivalent to the existence of an absent-minded agent.*

- (a) *There exist \bar{h}, \bar{t} , and $0 \leq m < n \leq N(\bar{t})$ such that $\{ {}_1\bar{t}_m, {}_1\bar{t}_n \} \subseteq \bar{h}$.*
- (b) *There exist \bar{h}, \bar{t} , and $0 \leq m < N(\bar{t})$ such that $\{ {}_1\bar{t}_m, \bar{t} \} \subseteq \bar{h}$.*
- (c) *There exist \bar{h}, \bar{t} , and $1 \leq m \leq N(\bar{t})$ such that $\bar{t}_m \in \bar{F}(\bar{h})$ and $\bar{t} \in \bar{h}$.*
- (d) *There exist \bar{t} and $1 \leq m < n \leq N(\bar{t})$ such that $\bar{t}_m = \bar{t}_n$.*
- (e) *There exist \bar{h}, \bar{t} , and $1 \leq m < n \leq N(\bar{t})$ such that $\{ \bar{t}_m, \bar{t}_n \} \subseteq \bar{F}(\bar{h})$.*

Proof. By inspection, (a) is equivalent to the existence of an absent-minded agent.

(a) \Rightarrow (b). If (a) holds for $\bar{t} = \bar{t}^*$ and $n = n^*$, then (b) holds for $\bar{t} = {}_1\bar{t}_{n^*}^*$.

(b) \Rightarrow (c). If (b) holds for $m = m^*$, then (c) holds for $m = m^* + 1$.

(c) \Rightarrow (d). Assume (c). Since $\bar{t}_m \in \bar{F}(\bar{h})$ and $\bar{t} \in \bar{h}$, it must be that $\bar{t}^* = \bar{t} \oplus (\bar{t}_m)$ belongs to \bar{T} . Thus (d) holds at $\bar{t} = \bar{t}^*$ because both \bar{t}_m^* and $\bar{t}_{N(\bar{t}^*)}^*$ equal \bar{t}_m .

(d) \Rightarrow (e) Assume (d). Since \bar{H} partitions $\bar{T} \sim \bar{Z}$, there is an \bar{h} such that ${}_1\bar{t}_{m=1} \in \bar{h}$ and hence $\bar{t}_m \in \bar{F}(\bar{h})$. Since $\bar{F}(\bar{h})$ has \bar{t}_m as an element, it must have the singleton $\{ \bar{t}_m, \bar{t}_n \}$ as a subset. Thus (e) holds.

(e) \Rightarrow (a). If (e) holds at $m = m^*$ and $n = n^*$, then (a) holds at $m = m^* - 1$ and $n = n^* - 1$. \square

Lemma A.5. *Agent recall is equivalent to $(\forall \bar{t}) |R(\bar{t})| = N(\bar{t})$.*

Proof. By Lemma A.4(d), the negation of agent recall is equivalent to the existence of a \bar{t} such that $|R(\bar{t})| < N(\bar{t})$. This is equivalent to

the negation of $(\forall \bar{t}) |R(\bar{t})| = N(\bar{t})$ since $|R(\bar{t})|$ can never exceed $N(\bar{t})$.
 \square

A.4. REDUCING SEQUENCES TO SETS

Lemma A.6 (“The Zipper”).¹¹ *If (A, \bar{T}) has agent recall, then*

$$(\forall \bar{t}, \bar{t}^*) R(\bar{t}) \supseteq R(\bar{t}^*) \Rightarrow {}_1\bar{t}_{N(\bar{t}^*)} = \bar{t}^* .$$

Proof. Take any \bar{t} and \bar{t}^* such that $R(\bar{t}) \supseteq R(\bar{t}^*)$. By Lemma A.5, by $R(\bar{t}) \supseteq R(\bar{t}^*)$, and by Lemma A.5 again, we have

$$N(\bar{t}) = |R(\bar{t})| \geq |R(\bar{t}^*)| = N(\bar{t}^*) .$$

The next two paragraphs will show by induction on $n \in \{1, 2, \dots, N(\bar{t}^*)\}$ that $(\forall n \leq N(\bar{t}^*)) {}_1\bar{t}_n = {}_1\bar{t}_n^*$.

For the initial step at $n = 1$, suppose that $\bar{t}_1 \neq \bar{t}_1^*$. Let \bar{h} be the agent containing the initial node $\{\}$ and note that $\{\bar{t}_1, \bar{t}_1^*\} \subseteq F(\bar{h})$ (in fact, agent recall implies that \bar{h} must be $\{\{\}\}$ but this observation is superfluous here). Since $R(\bar{t}) \supseteq R(\bar{t}^*)$, it must be that $\bar{t}_1^* \in R(\bar{t})$, hence there exists a $k > 1$ such that $\bar{t}_k = \bar{t}_1^*$, and hence, by the previous sentence, there exists a $k > 1$ such that $\{\bar{t}_1, \bar{t}_k\} \subseteq \bar{F}(\bar{h})$. Thus by Lemma A.4(e) there is an absent-minded agent. This violates agent recall, and hence, it must be that $\bar{t}_1 = \bar{t}_1^*$.

For the inductive step at $n \in \{2, 3, \dots, N(\bar{t}^*)\}$, assume that ${}_1\bar{t}_{n-1} = {}_1\bar{t}_{n-1}^*$ and suppose that $\bar{t}_n \neq \bar{t}_n^*$. Let \bar{h} be the agent containing ${}_1\bar{t}_{n-1}$ ($= {}_1\bar{t}_{n-1}^*$) and note that $\{\bar{t}_n, \bar{t}_n^*\} \subseteq F(\bar{h})$. Since $R(\bar{t}) \supseteq R(\bar{t}^*)$, it must be that $\bar{t}_n^* \in R(\bar{t})$, hence there exists a $m \neq n$ such that $\bar{t}_m = \bar{t}_n^*$, and hence, by the previous sentence, there exists a $m \neq n$ such that $\{\bar{t}_n, \bar{t}_m\} \subseteq \bar{F}(\bar{h})$. Thus by Lemma A.4(e) there is an absent-minded agent. This violates agent recall, and hence, it must be that $\bar{t}_n = \bar{t}_n^*$.

Therefore $(\forall n \leq N(\bar{t}^*)) {}_1\bar{t}_n = {}_1\bar{t}_n^*$. In particular, at $n = N(\bar{t}^*)$, we have ${}_1\bar{t}_{N(\bar{t}^*)} = {}_1\bar{t}_{N(\bar{t}^*)}^*$. The right-hand side is \bar{t}^* . \square

Lemma A.7. *Every sequence-tree game with agent recall is isomorphic to exactly one set-tree game.*

Proof. Let $(A, \bar{T}, \bar{H}, \bar{I}, \bar{v}^c, \bar{\rho}, \bar{u})$ be a sequence-tree game with agent recall, and derive \bar{F} and \bar{Z} from (A, \bar{T}) .

¹¹The name refers to the lemma’s inductive proof, which starts with the sequences’ first actions and works its way up.

Step 1: Uniqueness. Suppose that both $(A, T, H, I, i^c, \rho, u)$ and $(A, T', H', I', (i^c)', \rho', u')$ are isomorphic to the given $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$. By (5a), we have $T = T'$. Further, by (6a,b,c,d), we have $(H, I, i^c, \rho) = (H', I', (i^c)', \rho')$. Finally, by (6b,c,e) and Lemma A.3(e), we have $u = u'$.

Step 2: Two preliminary observations. This paragraph shows

$$(14) \quad (\forall \bar{t}^*, a, \bar{t}) \quad \bar{t}^* \oplus(a) = \bar{t} \Rightarrow a \notin R(\bar{t}^*) \text{ and } R(\bar{t}^*) \cup \{a\} = R(\bar{t}) .$$

Accordingly, take any \bar{t}^* , a , and \bar{t} such that $\bar{t}^* \oplus(a) = \bar{t}$. Note

$$|R(\bar{t}^*)| + 1 = N(\bar{t}^*) + 1 = N(\bar{t}) = |R(\bar{t})| .$$

by Lemma A.5, by $\bar{t}^* \oplus(a) = \bar{t}$, and by Lemma A.5 again. This and $\bar{t}^* \oplus(a) = \bar{t}$ yield $a \notin R(\bar{t}^*)$, which is the first fact to be derived. Further, $\bar{t}^* \oplus(a) = \bar{t}$ also implies that $R(\bar{t}) = R(\bar{t}^* \oplus(a)) = R(\bar{t}^*) \cup \{a\}$, which is the second fact to be derived.

Conversely, this paragraph shows

$$(15) \quad (\forall \bar{t}^*, a, \bar{t}) \quad \bar{t}^* \oplus(a) = \bar{t} \Leftarrow a \notin R(\bar{t}^*) \text{ and } R(\bar{t}^*) \cup \{a\} = R(\bar{t}) .$$

Accordingly, take any \bar{t}^* , a , and \bar{t} such that $a \notin R(\bar{t}^*)$ and $R(\bar{t}^*) \cup \{a\} = R(\bar{t})$. Note

$$N(\bar{t}^*) + 1 = |R(\bar{t}^*)| + 1 = |R(\bar{t})| = N(\bar{t}) .$$

by Lemma A.5, by the assumption of the previous sentence, and by Lemma A.5 again. Since $R(\bar{t}) = R(\bar{t}^*) \cup \{a\} \supseteq R(\bar{t}^*)$, Lemma A.6 (the “zipper”) shows that ${}_1\bar{t}_{N(\bar{t}^*)} = \bar{t}^*$. Thus by the last two sentences together, ${}_1\bar{t}_{N(\bar{t})-1} = \bar{t}^*$. Therefore, since $\{a\} = R(\bar{t}) \sim R(\bar{t}^*)$, it must be that $\bar{t}_{N(\bar{t})} = a$. The last two sentences together yield $\bar{t} = \bar{t}^* \oplus(a)$.

Step 3: An isomorphic set tree. Define (A, T) by letting $T = R_1(\bar{T})$.

This paragraph shows

$$(16) \quad R|_{\bar{T}} \text{ is an invertible function from } \bar{T} \text{ onto } T .$$

Since $T = R_1(\bar{T})$ by definition, we only need show that $R|_{\bar{T}}$ is injective. Accordingly, suppose that \bar{t} and \bar{t}^* are elements of \bar{T} such that $R(\bar{t}) = R(\bar{t}^*)$. By Lemma A.6 (the “zipper”), we have ${}_1\bar{t}_{N(\bar{t}^*)} = \bar{t}^*$. Further, the left-hand side is \bar{t} because

$$N(\bar{t}^*) = |R(\bar{t}^*)| = |R(\bar{t})| = N(\bar{t})$$

by Lemma A.5, by $R(\bar{t}) = R(\bar{t}^*)$, and by Lemma A.5 again.

This paragraph shows that

$$(17) \quad (\forall t^*, a, t) \\ (R|_{\bar{T}})^{-1}(t^*) \oplus (a) = (R|_{\bar{T}})^{-1}(t) \Leftrightarrow a \notin t^* \text{ and } t^* \cup \{a\} = t .$$

Accordingly, take any t^* , a , and t . Then define $\bar{t}^* = (R|_{\bar{T}})^{-1}(t^*)$ and $\bar{t} = (R|_{\bar{T}})^{-1}(t)$ and note

$$\begin{aligned} & (R|_{\bar{T}})^{-1}(t^*) \oplus (a) = (R|_{\bar{T}})^{-1}(t) \\ \Leftrightarrow & \bar{t}^* \oplus (a) = \bar{t} \\ \Leftrightarrow & a \notin R(\bar{t}^*) \text{ and } R(\bar{t}^*) \cup \{a\} = R(\bar{t}) \\ \Leftrightarrow & a \notin t^* \text{ and } t^* \cup \{a\} = t , \end{aligned}$$

where the first equivalence follows from the definitions of \bar{t}^* and \bar{t} , the second from (14) and (15), and the third from the definitions of \bar{t}^* and \bar{t} and from the invertibility (16) of $R|_{\bar{T}}$.

We now show that (A, T) is a set tree. Since $A = \bigcup \{R(\bar{t})|\bar{t}\}$ by assumption and since $\{R(\bar{t})|\bar{t}\} = T$ by the definition of T , we have that $A = \bigcup T$. It remains to be shown that every nonempty t has a unique last action.

Accordingly, take any nonempty t . First consider uniqueness. By (17) in the direction \Leftarrow , every last action of t must be the last element of the sequence $(R|_{\bar{T}})^{-1}(t)$. To see existence, define $\bar{t} = (R|_{\bar{T}})^{-1}(t)$, and then from this \bar{t} derive $t^* = R({}_1\bar{t}_{N(\bar{t})-1})$ and $a = \bar{t}_{N(\bar{t})}$. Then by substitution and manipulation,

$$\begin{aligned} & (R|_{\bar{T}})^{-1}(t^*) \oplus (a) \\ &= (R|_{\bar{T}})^{-1}(R({}_1\bar{t}_{N(\bar{t})-1})) \oplus (\bar{t}_{N(\bar{t})}) \\ &= {}_1\bar{t}_{N(\bar{t})-1} \oplus (\bar{t}_{N(\bar{t})}) \\ &= \bar{t} \\ &= (R|_{\bar{T}})^{-1}(t) . \end{aligned}$$

Since this is the left-hand side of (17), we have the right-hand side of (17), which states that this a is a last action of t .

Finally, we note that (A, \bar{T}) and (A, T) are isomorphic by (16), (14), and (15).

Step 4: An isomorphic set-tree game. Derive F and Z from (A, T) . Then define (H, I, i^c, ρ, u) by

$$(18a) \quad H = \{ R_1(\bar{h}) \mid \bar{h} \in \bar{H} \}$$

$$(18b) \quad I = \{ R_2(\bar{i}) \mid \bar{i} \in \bar{I} \}$$

$$(18c) \quad i^c = R_2(\bar{i}^c)$$

$$(18d) \quad \rho = \bar{\rho} \text{ and}$$

$$(18e) \quad (\forall i \neq i^c)(\forall t \in Z) u_i(t) = \bar{u}_{(R_2|_{\mathcal{P}^2(\bar{T})})^{-1}(i)}((R|_{\bar{T}})^{-1}(t)) .$$

This paragraph derives assumption (4a). Accordingly, take any t^1 , t^2 , and h , and define $\bar{t}^1 = (R|_{\bar{T}})^{-1}(t^1)$, $\bar{t}^2 = (R|_{\bar{T}})^{-1}(t^2)$, and $\bar{h} = (R_1|_{\mathcal{P}(\bar{T})})^{-1}(h)$. Then

$$\begin{aligned} & \{t^1, t^2\} \subseteq h \\ \Rightarrow & \{\bar{t}^1, \bar{t}^2\} \subseteq \bar{h} \\ \Rightarrow & \bar{F}(\bar{t}^1) = \bar{F}(\bar{t}^2) \\ \Rightarrow & F(t^1) = F(t^2) , \end{aligned}$$

where the second implication follows from (2a) and from $\bar{h} \in \bar{H}$ by (18a), and the last implication follows from Lemma A.3(a).

We now derive assumption (4b). Accordingly, take any t^1 and t^2 , and define $\bar{t}^1 = (R|_{\bar{T}})^{-1}(t^1)$ and $\bar{t}^2 = (R|_{\bar{T}})^{-1}(t^2)$. Then

$$\begin{aligned} & F(t^1) \cap F(t^2) \neq \emptyset \\ \Rightarrow & \bar{F}(\bar{t}^1) \cap \bar{F}(\bar{t}^2) \neq \emptyset \\ \Rightarrow & (\exists \bar{h}) \{\bar{t}^1, \bar{t}^2\} \subseteq \bar{h} \\ \Rightarrow & (\exists h) \{t^1, t^2\} \subseteq h , \end{aligned}$$

where the first implication follows from Lemma A.3(a), the second from (2b), and the last from (18a).

By Step 3 in this proof, (A, T) is a set tree. Further, by Lemma A.3(d) and (18a), and by the last two paragraphs, H is a partition of $T \sim Z$ that satisfies property (4). Hence $(A, T, H, I, i^c, \rho, u)$ is a set-tree game.

Finally, by Step 3 in this proof, (A, T) and (A, \bar{T}) are isomorphic. Additionally, (18) and Lemma A.3(e) imply (6). Hence $(A, T, H, I, i^c, \rho, u)$ and $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$ are isomorphic. \square

A.5. CONSTRUCTING SEQUENCES FROM SETS

Lemma A.8. *Every set-tree game is isomorphic to exactly one sequence-tree game, and that sequence-tree game has agent recall.*

Proof. Let $(A, T, H, I, i^c, \rho, u)$ be a set-tree game, derive F and Z from (A, T) , and let $\alpha_*: T \rightarrow A$ be the function that takes each node $t \in T$ to its unique last action.

Step 1: Uniqueness. Suppose that $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$ and $(A, \bar{\bar{T}}, \bar{\bar{H}}, \bar{\bar{I}}, \bar{\bar{i}}^c, \bar{\bar{\rho}}, \bar{\bar{u}})$ are two sequence-tree games that are isomorphic to $(A, T, H, I, i^c, \rho, u)$.

This and the next two paragraphs show that $\bar{T} = \bar{\bar{T}}$. Suppose not. Then because both (A, \bar{T}) and $(A, \bar{\bar{T}})$ satisfy isomorphism condition (5a), there must be $\bar{t}, \bar{\bar{t}}$, and t such that $\bar{t} \neq \bar{\bar{t}}$ and yet $R(\bar{t}) = R(\bar{\bar{t}}) = t$.

This long paragraph shows by induction that

$$\begin{aligned} & (\forall k \in \{0, 1, \dots, |t|\}) \\ (19a) \quad & {}_1\bar{t}_{N(\bar{t})-k} \neq {}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k} , \\ (19b) \quad & R({}_1\bar{t}_{N(\bar{t})-k}) = R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k}) , \\ (19c) \quad & \text{and } |R({}_1\bar{t}_{N(\bar{t})-k})| = |t| - k . \end{aligned}$$

The initial step at $k=0$ follows from the definition of \bar{t} , $\bar{\bar{t}}$, and t . Now assume that (19) holds at $k < |t|$. By the definitions of \bar{t} , $\bar{\bar{t}}$, and t , it must be $N(\bar{t})$ and $N(\bar{\bar{t}})$ are at least as big as $|t|$ and thus strictly bigger than k . As a result, we may write

$$(20) \quad \begin{aligned} & {}_1\bar{t}_{N(\bar{t})-k-1} \oplus (\bar{t}_{N(\bar{t})-k}) = {}_1\bar{t}_{N(\bar{t})-k} \quad \text{and} \\ & {}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1} \oplus (\bar{\bar{t}}_{N(\bar{\bar{t}})-k}) = {}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k} . \end{aligned}$$

Thus, by applying isomorphism property (5b) twice, we find

$$(21) \quad \begin{aligned} & \bar{t}_{N(\bar{t})-k} \notin R({}_1\bar{t}_{N(\bar{t})-k-1}) , \\ & R({}_1\bar{t}_{N(\bar{t})-k-1}) \cup \{\bar{t}_{N(\bar{t})-k}\} = R({}_1\bar{t}_{N(\bar{t})-k}) , \\ & \text{and } \bar{\bar{t}}_{N(\bar{\bar{t}})-k} \notin R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1}) , \\ & R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1}) \cup \{\bar{\bar{t}}_{N(\bar{\bar{t}})-k}\} = R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k}) . \end{aligned}$$

Thus, by applying the definition of last action twice, we find

$$\begin{aligned} \bar{t}_{N(\bar{t})-k} &= \alpha_*(R({}_1\bar{t}_{N(\bar{t})-k})) \quad \text{and} \\ \bar{\bar{t}}_{N(\bar{\bar{t}})-k} &= \alpha_*(R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k})) . \end{aligned}$$

But by (19b), the right-hand sides of these two equalities must be equal. Thus we may define a_* to be equal to both $\bar{t}_{N(\bar{t})-k}$ and $\bar{\bar{t}}_{N(\bar{\bar{t}})-k}$, and then

substitute out both of these latter terms in (20) and (21) to obtain

$$(22) \quad \begin{aligned} {}_1\bar{t}_{N(\bar{t})-k-1} \oplus (a_*) &= {}_1\bar{t}_{N(\bar{t})-k} \text{ and} \\ {}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1} \oplus (a_*) &= {}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k} . \end{aligned}$$

and

$$(23) \quad \begin{aligned} a_* \notin R({}_1\bar{t}_{N(\bar{t})-k-1}) \text{ and } R({}_1\bar{t}_{N(\bar{t})-k-1}) \cup \{a_*\} &= R({}_1\bar{t}_{N(\bar{t})-k}) \text{ and} \\ a_* \notin R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1}) \text{ and } R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1}) \cup \{a_*\} &= R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k}) . \end{aligned}$$

By (19a), the pair (22) implies that

$${}_1\bar{t}_{N(\bar{t})-k-1} \neq {}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1} .$$

The pair (23) implies that

$$\begin{aligned} R({}_1\bar{t}_{N(\bar{t})-k-1}) &= R({}_1\bar{t}_{N(\bar{t})-k}) \sim \{a_*\} \text{ and} \\ R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1}) &= R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k}) \sim \{a_*\} , \end{aligned}$$

and thus by (19b) we have that

$$R({}_1\bar{t}_{N(\bar{t})-k-1}) = R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-k-1}) .$$

Finally, the first half of (23) together with (19c) imply that

$$|R({}_1\bar{t}_{N(\bar{t})-k-1})| = |R({}_1\bar{t}_{N(\bar{t})-k})| - 1 = |t| - k - 1 .$$

The last three sentences have derived (19) at $k+1$.

At $k = |t|$, equations (19b) and (19c) imply that both $R({}_1\bar{t}_{N(\bar{t})-|t|})$ and $R({}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-|t|})$ are empty. Thus both ${}_1\bar{t}_{N(\bar{t})-|t|}$ and ${}_1\bar{\bar{t}}_{N(\bar{\bar{t}})-|t|}$ are empty, in contradiction to (19a). Therefore $\bar{T} = \bar{\bar{T}}$.

Next, we show $(\bar{H}, \bar{I}, \bar{i}^c) = (\bar{\bar{H}}, \bar{\bar{I}}, \bar{\bar{i}}^c)$. Since $\bar{T} = \bar{\bar{T}}$, we have that $R|_{\bar{T}} = R|_{\bar{\bar{T}}}$, that $R_1|_{\mathcal{P}(\bar{T})} = R_1|_{\mathcal{P}(\bar{\bar{T}})}$, and that $R_2|_{\mathcal{P}^2(\bar{T})} = R_2|_{\mathcal{P}^2(\bar{\bar{T}})}$. Therefore, since both \bar{H} and $\bar{\bar{H}}$ satisfy (6a), we have

$$(24a) \quad \bar{H} = \{(R_1|_{\mathcal{P}(\bar{T})})^{-1}(h) | h \in H\} = \{(R_1|_{\mathcal{P}(\bar{\bar{T}})})^{-1}(h) | h \in H\} = \bar{\bar{H}} ,$$

and since both \bar{I} and $\bar{\bar{I}}$ satisfy (6b), we have

$$(24b) \quad \bar{I} = \{(R_2|_{\mathcal{P}^2(\bar{T})})^{-1}(i) | i \in I\} = \{(R_2|_{\mathcal{P}^2(\bar{\bar{T}})})^{-1}(i) | i \in I\} = \bar{\bar{I}} ,$$

and since both \bar{i}^c and $\bar{\bar{i}}^c$ satisfy (6c), we have

$$(24c) \quad \bar{i}^c = (R_2|_{\mathcal{P}^2(\bar{T})})^{-1}(i^c) = (R_2|_{\mathcal{P}^2(\bar{\bar{T}})})^{-1}(i^c) = \bar{\bar{i}}^c .$$

Finally, we show $(\bar{\rho}, \bar{u}) = (\bar{\bar{\rho}}, \bar{\bar{u}})$. Trivially, $\bar{\rho} = \rho = \bar{\bar{\rho}}$ since both $\bar{\rho}$ and $\bar{\bar{\rho}}$ satisfy (6d). To get at the payoff functions, begin by deriving \bar{Z} from (A, \bar{T}) and $\bar{\bar{Z}}$ from $(A, \bar{\bar{T}})$. Then since $\bar{I} \sim \{\bar{i}^c\} = \bar{\bar{I}} \sim \{\bar{\bar{i}}^c\}$ by (24b,c), and since $\bar{Z} = \bar{\bar{Z}}$ because $\bar{T} = \bar{\bar{T}}$, we have that $(\bar{I} \sim \{\bar{i}^c\}) \times \bar{Z} = (\bar{\bar{I}} \sim \{\bar{\bar{i}}^c\}) \times \bar{\bar{Z}}$,

or in other words, that the domain of \bar{u} equals the domain of $\bar{\bar{u}}$. Then, for any (\bar{i}, \bar{t}) in that common domain, we have

$$\bar{u}_{\bar{i}}(\bar{t}) = u_{R_2(\bar{i})}(R(\bar{t})) = \bar{\bar{u}}_{\bar{i}}(\bar{t})$$

because both \bar{u} and $\bar{\bar{u}}$ satisfy (6e) (the single bars on \bar{i} and \bar{t} on the right-hand side are correct). The last two sentences imply $\bar{u} = \bar{\bar{u}}$.

Step 2: Define \bar{T} . We now begin the task of constructing a sequence-tree game which is isomorphic to $(A, T, H, I, i^c, \rho, u)$. The first job is to define \bar{T} .

For any $n \geq 0$, let $T_n = \{ t \mid |t| = n \}$ be the set of nodes with n elements. Because A is finite, there is some N such that $T_N \neq \emptyset$ and $(\forall n > N) T_n = \emptyset$. Thus $T = \bigcup_{n=0}^N T_n$. Further, let t^N be some element of T_N , and for all $n \in \{0, 1, 2, \dots, N-1\}$, let t^n be $t^{n+1} \sim \{\alpha_*(t^{n+1})\}$. Since each $t^n \in T_n$, we have shown that $(\forall n \leq N) T_n \neq \emptyset$. In particular, $T_0 \neq \emptyset$ and thus $T_0 = \{\{\}\}$.

We now define a sequence $\langle Q_n \rangle_{n=0}^N$ of functions in which each function Q_n maps each node t of T_n to some finite action sequence \bar{t} . We do this recursively. To begin, recall $T_0 = \{\{\}\}$ from the previous paragraph and define the one-element function Q_0 by $Q_0(\{\}) = \{\}$. Thus the empty set $t = \{\}$ is mapped to the empty sequence $\bar{t} = \{\}$. Then, for any $n \in \{1, 2, \dots, N\}$, use Q_{n-1} to define Q_n at each $t \in T_n$ by

$$(25) \quad Q_n(t) = Q_{n-1}(t \sim \{\alpha_*(t)\}) \oplus (\alpha_*(t)) .$$

Note that $Q_{n-1}(t \sim \{\alpha_*(t)\})$ is well-defined because $t \sim \{\alpha_*(t)\}$ has $n-1$ elements because $t \in T_n$ and $\alpha_*(t)$ is its last action.

Define $\bar{T} = \bigcup_n Q_n(T_n)$, where here, and for the remainder of the proof, we implicitly assume that n ranges over $\{0, 1, \dots, N\}$.

Step 3: Show (A, \bar{T}) is a sequence tree. First we show by induction that

$$(26) \quad (\forall n)(\forall t \in T_n) R(Q_n(t)) = t .$$

This holds at $n=0$ because $R(Q_0(\{\})) = R(\{\}) = \{\}$. Further, it holds at $n \geq 1$ if it holds at $n-1$ because

$$\begin{aligned} (\forall t \in T_n) R(Q_n(t)) &= R(Q_{n-1}(t \sim \{\alpha_*(t)\}) \oplus (\alpha_*(t))) \\ &= R(Q_{n-1}(t \sim \{\alpha_*(t)\})) \cup R((\alpha_*(t))) \\ &= t \sim \{\alpha_*(t)\} \cup \{\alpha_*(t)\} \end{aligned}$$

$$= t ,$$

where the first equality holds by the definition (25) of Q_n , and the third holds by the inductive hypothesis.

This observation allows us to claim

$$(27) \quad A = \bigcup_{\bar{t}} R(\bar{t}) .$$

Easily, $A \supseteq \bigcup_{\bar{t}} R(\bar{t})$ because each $R(\bar{t})$ is a set of actions. Conversely, take any a . By assumption there is some t such that $a \in t$. Then by construction there is some n such that $t \in T_n$. Thus by (26), we have $a \in t = R(Q_n(t))$. Therefore, since $Q_n(t) \in Q_n(T_n) \subseteq \bar{T}$, this $Q_n(t)$ is a \bar{t} such that $a \in R(\bar{t})$.

Next we show by induction that

$$(28) \quad (\forall n)(\forall t \in T_n) N(Q_n(t)) = n .$$

This holds at $n = 0$ because $N(Q_0(\{\})) = N(\{\}) = 0$. Further, it holds at any $n \geq 1$ if it holds at $n-1$ because

$$\begin{aligned} (\forall t \in T_n) N(Q_n(t)) &= N(Q_{n-1}(t \sim \{\alpha_*(t)\}) \oplus (\alpha_*(t))) \\ &= N(Q_{n-1}(t \sim \{\alpha_*(t)\})) + N((\alpha_*(t))) \\ &= (n-1) + 1 \\ &= n , \end{aligned}$$

where the first equality holds by the definition (25) of Q_n , and the third by the inductive hypothesis.

This observation allows us to claim

$$(29) \quad (\forall n) \{ \bar{t} \in \bar{T} \mid N(\bar{t})=n \} = Q_n(T_n) .$$

The inclusion \supseteq follows from (28) at n . Conversely, if there were an element of $\{ \bar{t} \in \bar{T} \mid N(\bar{t})=n \}$ that was from $Q_m\{T_m\}$ for some $m \neq n$ it would violate (28) at m .

Finally we show that \bar{T} satisfies assumption (1). Accordingly, take any $\bar{t} \in \bar{T}$. By (29), there exists $t \in T_{N(\bar{t})}$ such that $\bar{t} = Q_{N(\bar{t})}(t)$, and thus the definition (25) of $Q_{N(\bar{t})}$ yields that

$${}_1\bar{t}_{N(\bar{t})-1} = Q_{N(\bar{t})-1}(t \sim \{a_*(t)\}) \in \bar{T} .$$

By (27) and the previous paragraph, (A, \bar{T}) is a sequence tree.

Step 4: Show isomorphism between trees. Next we show that the sequence tree (A, \bar{T}) is isomorphic to the original set tree (A, T) . In

particular, the next paragraph shows $R|_{\bar{T}}$ is invertible and the remainder of this step shows that concatenation is isomorphic to union.

Define $Q = \bigcup_n Q_n$. The remainder of this paragraph shows (32) below. To begin, (26) implies that each $R|_{Q_n(T_n)}$ is the inverse of Q_n . In other words,

$$(30) \quad (\forall n) \quad Q_n = (R|_{Q_n(T_n)})^{-1} \text{ is} \\ \text{an invertible function from } T_n \text{ onto } Q_n(T_n) .$$

The domain of Q is $T = \bigcup_n T_n$ and its range is $\bar{T} = \bigcup_n Q_n(T_n)$. Further, T is partitioned by $\{T_n\}_n$, and \bar{T} is partitioned by $\{Q_n(T_n)\}_n$ (because of (29)). Therefore (30) implies that

$$(31) \quad Q = (R|_{\bar{T}})^{-1} \text{ is an invertible function from } T \text{ onto } \bar{T} .$$

This is equivalent to

$$(32) \quad R|_{\bar{T}} = Q^{-1} \text{ is an invertible function from } \bar{T} \text{ onto } T .$$

To begin proving that concatenation is isomorphic to union, this paragraph shows

$$(33) \quad (\forall n \geq 1)(\forall t^* \in T_{n-1})(\forall a)(\forall t \in T_n) \\ Q_{n-1}(t^*) \oplus (a) = Q_n(t) \Leftrightarrow a \notin t^* \text{ and } t^* \cup \{a\} = t .$$

Accordingly, take any such n , t^* , a , and t . Then

$$\begin{aligned} Q_{n-1}(t^*) \oplus (a) &= Q_n(t) \\ \Leftrightarrow Q_{n-1}(t^*) \oplus (a) &= Q_{n-1}(t \sim \{\alpha_*(t)\}) \oplus (\alpha_*(t)) \\ \Leftrightarrow Q_{n-1}(t^*) &= Q_{n-1}(t \sim \{\alpha_*(t)\}) \text{ and } a = \alpha_*(t) \\ \Leftrightarrow t^* &= t \sim \{\alpha_*(t)\} \text{ and } a = \alpha_*(t) \\ \Leftrightarrow a \notin t^* &\text{ and } t^* \cup \{a\} = t \end{aligned}$$

where the first equivalence holds by the definition of Q_n at (25), the second equivalence by breaking the vector equality into two components, the third equivalence by applying R and (32) to the first equality, and the fourth equivalence by $\alpha_*(t)$ being a last action.

Essentially, this paragraph removes the n from (33). Specifically, it shows that

$$(34) \quad (\forall t^*)(\forall a)(\forall t) \\ Q(t^*) \oplus (a) = Q(t) \Leftrightarrow a \notin t^* \text{ and } t^* \cup \{a\} = t .$$

First suppose t^* , a , and t satisfy $Q(t^*) \oplus (a) = Q(t)$ and let $n = |t|$. By (28) and the definition of Q , we have $Q(t) = Q_n(t)$ and $Q(t^*) = Q_{n-1}(t^*)$. Hence $a \notin t^*$ and $t^* \cup \{a\} = t$ by (33). Conversely, suppose t^* , a , and t satisfy $a \notin t^*$ and $t^* \cup \{a\} = t$ and let $n = |t|$. Then $n-1 = |t^*|$. Thus since $t \in T_n$ and $t^* \in T_{n-1}$, (33) yields that $Q_{n-1}(t^*) \oplus (a) = Q_n(t)$. By the definition of Q , this is equivalent to $Q(t^*) \oplus (a) = Q(t)$.

Essentially, this final paragraph quantifies (34) in terms of sequences rather than sets. Specifically, it shows that

$$(35) \quad (\forall \bar{t}^*, a, \bar{t}) \\ \bar{t}^* \oplus (a) = \bar{t} \Leftrightarrow a \notin R(\bar{t}^*) \text{ and } R(\bar{t}^*) \cup \{a\} = R(\bar{t}) .$$

Accordingly, take any \bar{t}^* , a , and \bar{t} , define $t^* = R(\bar{t}^*)$, and define $t = R(\bar{t})$. Then

$$\begin{aligned} \bar{t}^* \oplus (a) &= \bar{t} \\ \Leftrightarrow Q(t^*) \oplus (a) &= Q(t) \\ \Leftrightarrow a \notin t^* \text{ and } t^* \cup \{a\} &= t \\ \Leftrightarrow a \notin R(\bar{t}^*) \text{ and } R(\bar{t}^*) \cup \{a\} &= R(\bar{t}) , \end{aligned}$$

where the first equivalence holds by the definitions of t^* and t and by the fact that $R|_{\bar{t}} = Q^{-1}$ by (32). The second equivalence holds by (34), and the third by the definitions of t^* and t .

Equations (32) and (35) show that the set tree (A, T) and the sequence tree (A, \bar{T}) are isomorphic.

Step 5: Define the sequence-tree game. Derive \bar{F} and \bar{Z} from (A, \bar{T}) . Then define $(\bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$ by

$$(36a) \quad \bar{H} = \{ (R_1|_{\mathcal{P}(\bar{T})})^{-1}(h) \mid h \in H \}$$

$$(36b) \quad \bar{I} = \{ (R_2|_{\mathcal{P}^2(\bar{T})})^{-1}(i) \mid i \in I \}$$

$$(36c) \quad \bar{i}^c = (R_2|_{\mathcal{P}^2(\bar{T})})^{-1}(i^c)$$

$$(36d) \quad \bar{\rho} = \rho \text{ and}$$

$$(36e) \quad (\forall \bar{i} \neq \bar{i}^c)(\forall \bar{t} \in \bar{Z}) \quad \bar{u}_{\bar{i}}(\bar{t}) = u_{R_2(\bar{i})}(R(\bar{t})) .$$

Since $R_1|_{\mathcal{P}(\bar{T})}$ and $R_2|_{\mathcal{P}^2(\bar{T})}$ are invertible because R is invertible, equations (36a,b,c) are equivalent to

$$(37a) \quad H = \{ R_1(\bar{h}) \mid \bar{h} \in \bar{H} \}$$

$$(37b) \quad I = \{ R_2(\bar{i}) \mid \bar{i} \in \bar{I} \}$$

$$(37c) \quad \text{and } i^c = R_2(\bar{i}^c) .$$

This paragraph derives assumption (2a). Accordingly, take any \bar{t}^1 , \bar{t}^2 , and \bar{h} . Then

$$\begin{aligned} & \{\bar{t}^1, \bar{t}^2\} \subseteq \bar{h} \\ \Rightarrow & \{R(\bar{t}^1), R(\bar{t}^2)\} \subseteq R_1(\bar{h}) \\ \Rightarrow & F(R(\bar{t}^1)) = F(R(\bar{t}^2)) \\ \Rightarrow & \bar{F}(\bar{t}^1) = \bar{F}(\bar{t}^2) , \end{aligned}$$

where the first implication follows from the definitions of R and R_1 , the second implication follows from assumption (4a) and the fact that $R_1(\bar{h}) \in H$ by (37a), and the last implication comes from Lemma A.3(a).

Then we derive assumption (2b). Accordingly, take any \bar{t}^1 and \bar{t}^2 . Then

$$\begin{aligned} & \bar{F}(\bar{t}^1) = \bar{F}(\bar{t}^2) \\ \Rightarrow & F(R(\bar{t}^1)) = F(R(\bar{t}^2)) \\ \Rightarrow & (\exists h)\{R(\bar{t}^1), R(\bar{t}^2)\} \subseteq h \\ \Rightarrow & (\exists h)\{\bar{t}^1, \bar{t}^2\} \subseteq (R_1|_{\mathcal{P}(\bar{T})})^{-1}(h) \\ \Rightarrow & (\exists \bar{h})\{\bar{t}^1, \bar{t}^2\} \subseteq \bar{h} , \end{aligned}$$

where the first implication follows from Lemma A.3(a), the second follows from assumption (4b), the third follows from the invertibility of $R|_{\bar{T}}$ and $R_1|_{\mathcal{P}(\bar{T})}$, and the fourth follows from the fact that $(R|_{\mathcal{P}(\bar{T})})^{-1}(h) \in \bar{H}$ by (36a).

By (37a) and Lemma A.3(d), and by the last two paragraphs, \bar{H} is a partition of $\bar{T} \sim \bar{Z}$ that satisfies assumption (2). Hence $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$ is a sequence-tree game.

Step 6: Show isomorphism between games. The trees (A, T) and (A, \bar{T}) are isomorphic by Step 4. Additionally, (37a,b,c) and (36d,e) imply (6). Hence $(A, T, H, I, i^c, \rho, u)$ is isomorphic to $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \bar{\rho}, \bar{u})$.

Step 7: Show agent recall. This last step could have been taken at any point after (31). Equation (26), the definition of T_n , and equation (28) yield that

$$(\forall n)(\forall t \in T_n) |R(Q_n(t))| = |t| = n = N(Q_n(t)) .$$

Thus by the definition of Q ,

$$(\forall t) |R(Q(t))| = N(Q(t)) .$$

Since Q is an invertible function from T onto \bar{T} by (31), this is equivalent to $(\forall \bar{t}) |R(\bar{t})| = N(\bar{t})$, which by Lemma A.5 is equivalent to agent recall. \square

APPENDIX B. THEOREM B AND ITS COROLLARIES

B.1. BASIC PROPERTIES OF A PLAUSIBILITY RELATION \succsim

Proof B.1 (for Lemma 3.1). Note $\overset{c}{\approx}$ is symmetric and equal to

$$(38) \quad \begin{aligned} & \{ (t, t \cup \{a\}) \mid a \in F(t) \text{ and } a \in A^c \} \\ & \cup \{ (t \cup \{a\}, t) \mid a \in F(t) \text{ and } a \in A^c \} . \end{aligned}$$

Further, $\overset{\sigma}{\approx}$ is symmetric, $\overset{\sigma}{\succ}$ is asymmetric, and the two are disjoint subsets of

$$(39) \quad \begin{aligned} & \{ (t, t \cup \{a\}) \mid a \in F(t) \text{ and } a \in A^s \} \\ & \cup \{ (t \cup \{a\}, t) \mid a \in F(t) \text{ and } a \in A^s \} . \end{aligned}$$

Similarly, $\overset{\beta}{\approx}$ is symmetric, $\overset{\beta}{\succ}$ is asymmetric, and the two are disjoint subsets of

$$(40) \quad \{ (t_1, t_2) \mid (\exists h \in H^s) \{t_1, t_2\} \in h \} .$$

This paragraph observes that these three sets are pairwise disjoint. (38) and (39) are disjoint because A is partitioned by $\{A^c, A^s\}$. Further, the union of (38) and (39) is disjoint from (40). If this were not the case, there would be t , a , and h such that $a \in F(t)$ and $\{t, t \cup \{a\}\} \in h$. Since $a \in F(t)$ and t and $t \cup \{a\}$ share an agent, assumption (4b) would imply that $a \in F(t \cup \{a\})$. However, this would contradict the definition of F , which would require that $a \notin t \cup \{a\}$.

Since the sets (38), (39), and (40) are pairwise disjoint, the disjointness observed in the first paragraph implies that \succsim is partitioned by $\{\overset{c}{\approx}, \overset{\sigma}{\approx}, \overset{\sigma}{\succ}, \overset{\beta}{\approx}, \overset{\beta}{\succ}\}$. Thus the symmetries and asymmetries observed in the first paragraph imply that \approx is partitioned by $\{\overset{c}{\approx}, \overset{\sigma}{\approx}, \overset{\beta}{\approx}\}$ and that \succ is partitioned by $\{\overset{\sigma}{\succ}, \overset{\beta}{\succ}\}$. \square

B.2. LEMMA ABOUT BINARY RELATIONS COMPARING SETS

Before proving Theorem B, we use Farkas' Lemma to derive a lemma about binary relations that compare sets. Although we could not find this exact result in the literature, it is one of many results closely related to Kraft, Pratt, and Seidenberg (1959) and Scott (1964). These papers are discussed in Subsection 3.3.

Fact B.2 (Farkas Lemma for Rational Matrices). *Let $D \in \mathbb{Q}^{dp}$ and $E \in \mathbb{Q}^{ep}$ be two rational matrices. Then the following are equivalent. ($D\pi \gg 0$ means every element of $D\pi$ is positive and δ^T means the transpose of δ .)*

(a) $(\exists \pi \in \mathbb{Z}^p)$ $D\pi \gg 0$ and $E\pi = 0$.

(b) Not $(\exists \delta \in \mathbb{Z}_+^d \sim \{0\})(\exists \varepsilon \in \mathbb{Z}^e)$ $\delta^T D + \varepsilon^T E = 0$.

(This fact is taken from Krantz, Luce, Suppes, and Tversky (1971, pages 62–63). We have replaced their $[\alpha_i]_{i=1}^{m'}$ with D and their $[\beta_i]_{i=1}^{m''}$ with E .)

Lemma B.3. *Let A be a finite set and let $\mathcal{P}(A)$ be the collection of all its subsets. Further, let \succsim be a binary relation over $\mathcal{P}(A)$, and let \succ and \approx be its asymmetric and symmetric parts. Then the following are equivalent (s and t denote arbitrary subsets of A).*

(a) (Density-Function Representation) *There exists $\pi: A \rightarrow \mathbb{Z}$ such that for all s and t*

$$\begin{aligned} s \succ t &\Rightarrow \sum_{a \in s} \pi(a) > \sum_{a \in t} \pi(a) \text{ and} \\ s \approx t &\Rightarrow \sum_{a \in s} \pi(a) = \sum_{a \in t} \pi(a) . \end{aligned}$$

(b) (Cancellation Laws) *There is no pair (s^m, t^m) from \succ in a finite sequence $\langle (s^m, t^m) \rangle_{m=1}^M$ of pairs from \succsim whenever the sequence satisfies $(\forall a) |\{m | a \in s^m\}| = |\{m | a \in t^m\}|$.*

Proof. (a \Rightarrow b) This is straightforward. See the paragraph containing equation (9).

(a \Leftarrow b) We begin by deriving two matrices from the relation \succsim . For any t , define the row vector $1^t \in \{0, 1\}^{|A|}$ by $1_a^t = 1$ if $a \in t$ and $1_a^t = 0$ if $a \notin t$. Then define the matrices $D = [1^s - 1^t]_{s \succ t}$ and $E = [1^s - 1^t]_{s \approx t}$ whose rows are indexed by the pairs of the relations \succ and \approx .

Now assume (b). This paragraph will argue that there cannot be column vectors $\delta \in \mathbb{Z}_+^{|\succ|} \sim \{0\}$ and $\varepsilon \in \mathbb{Z}^{|\approx|}$ such that $\delta^T D + \varepsilon^T E = 0$. To see this, suppose that there were such δ and ε . By the symmetry

of \approx , we may define $\varepsilon_+ \in \mathbb{Z}_+^{|\approx|}$ by

$$(\forall s \approx t) (\varepsilon_+)_{(s,t)} = \begin{pmatrix} \varepsilon_{(s,t)} - \varepsilon_{(t,s)} & \text{if } \varepsilon_{(s,t)} - \varepsilon_{(t,s)} \geq 0 \\ 0 & \text{otherwise} \end{pmatrix}$$

so that $\varepsilon^T E = \varepsilon_+^T E$. Thus we have $\delta \in \mathbb{Z}_+^{|\succ|} \sim \{0\}$ and $\varepsilon_+ \in \mathbb{Z}_+^{|\approx|}$ such that $\delta^T D + \varepsilon_+^T E = 0$. Now define the sequence $\langle (s^m, t^m) \rangle_{m=1}^M$ of pairs from \succ in such a way that every pair from \succ appears $\lambda_{(s,t)}$ times and every pair from \approx appears $(\mu_+)_{(s,t)}$ times. The equality $\delta^T D + \varepsilon_+^T E = 0$ yields that this sequence satisfies $(\forall a) |\{m | a \in s^m\}| = |\{m | a \in t^m\}|$, and $\delta \in \mathbb{Z}_+^{|\succ|} \sim \{0\}$ yields that it contains at least one pair from \succ . By condition (b), this is impossible.

Since the result of the previous paragraph is equivalent to condition (b) of Lemma B.2 (Farkas), there is a vector $\pi \in \mathbb{Z}^{|A|}$ such that $D\pi \gg 0$ and $E\pi = 0$. By the definitions of D and E , this is equivalent to condition (a) of this lemma. \square

B.3. PROOF OF THEOREM B

Lemma B.4 assembles elementary observations about consistency. Proof B.5 then combines Lemmas B.3 and B.4 to prove Theorem B.

Lemma B.4. *Take an assessment (σ, β) and its plausibility relation \succcurlyeq . Assume (σ, β) is consistent and let $\langle (\sigma_n, \beta_n) \rangle_n$ be a sequence of full-support Bayesian assessments that converges to it. Then*

$$\begin{aligned} (\forall t^1 \succ t^2) \lim_n \frac{\Pi_{a \in t^2}(\rho \cup \sigma_n)(a)}{\Pi_{a \in t^1}(\rho \cup \sigma_n)(a)} &= 0 \quad \text{and} \\ (\forall t^1 \approx t^2) \lim_n \frac{\Pi_{a \in t^2}(\rho \cup \sigma_n)(a)}{\Pi_{a \in t^1}(\rho \cup \sigma_n)(a)} &\in (0, \infty) \end{aligned}$$

(where $\Pi_{a \in \{t\}}(\rho \cup \sigma_n)(a)$ is defined to be one).

Proof. This paragraph shows

$$(41) \quad (\forall t^1 \succcurlyeq t^2) \lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} = 0 .$$

Accordingly, suppose $t^1 \succcurlyeq t^2$. By the definition of \succcurlyeq , there exists a such that $\sigma(a)=0$, $a \in F(t^1)$, and $t^1 \cup \{a\} = t^2$. Thus, since $a \notin t^1$ by the definition of F ,

$$\lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} = \lim_n \sigma_n(a) = \sigma(a) = 0 .$$

This paragraph shows

$$(42) \quad (\forall t^1 \stackrel{\sigma}{\approx} t^2) \lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} \in (0, \infty) .$$

Suppose $t^1 \stackrel{\sigma}{\approx} t^2$. By the definition of $\stackrel{\sigma}{\approx}$, there exists a with $\sigma(a) > 0$ such that either, $a \in F(t^1)$ and $t^1 \cup \{a\} = t^2$, or, $a \in F(t^2)$ and $t^2 \cup \{a\} = t^1$. In the first case, $a \notin t^1$ by the definition of F and thus

$$\lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} = \lim_n \sigma_n(a) = \sigma(a) \in (0, 1] ,$$

and in the second case, $a \notin t^2$ by the definition of F and thus

$$\lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} = \lim_n \frac{1}{\sigma_n(a)} = \frac{1}{\sigma(a)} \in [1, \infty) .$$

In a similar fashion, this paragraph shows

$$(43) \quad (\forall t^1 \stackrel{c}{\approx} t^2) \lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} \in (0, \infty) .$$

Suppose $t^1 \stackrel{c}{\approx} t^2$. By the definition of $\stackrel{c}{\approx}$, there exists $a \in A^c$ such that either, $a \in F(t^1)$ and $t^1 \cup \{a\} = t^2$, or, $a \in F(t^2)$ and $t^2 \cup \{a\} = t^1$. In the first case, $a \notin t^1$ and thus

$$\lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} = \lim_n \rho(a) \in (0, 1] ,$$

and in the second case, $a \notin t^2$ and thus

$$\lim_n \frac{\Pi_{a' \in t^2}(\rho \cup \sigma_n)(a')}{\Pi_{a' \in t^1}(\rho \cup \sigma_n)(a')} = \lim_n \frac{1}{\rho(a)} \in [1, \infty) .$$

Finally, note that if t^1 and t^2 share some $h \in H^s$, and if $\beta(t^1) > 0$, then

$$\begin{aligned} & \lim_n \frac{\Pi_{a \in t^2}(\rho \cup \sigma_n)(a)}{\Pi_{a \in t^1}(\rho \cup \sigma_n)(a)} \\ &= \lim_n \frac{\Pi_{a \in t^2}(\rho \cup \sigma_n)(a) / \sum_{t \in h} \Pi_{a \in t}(\rho \cup \sigma_n)(a)}{\Pi_{a \in t^1}(\rho \cup \sigma_n)(a) / \sum_{t \in h} \Pi_{a \in t}(\rho \cup \sigma_n)(a)} \\ &= \lim_n \frac{\beta_n(t^2)}{\beta_n(t^1)} = \frac{\beta(t^2)}{\beta(t^1)} , \end{aligned}$$

where the second equality follows from (7), and the third follows from consistency and the assumption that $\beta(t^1) > 0$. Thus by the definitions of $\stackrel{\beta}{\succ}$ and $\stackrel{\beta}{\approx}$ we have

$$(44) \quad (\forall t^1 \stackrel{\beta}{\succ} t^2) \lim_n \frac{\Pi_{a \in t^2}(\rho \cup \sigma_n)(a)}{\Pi_{a \in t^1}(\rho \cup \sigma_n)(a)} = \frac{\beta(t^2)}{\beta(t^1)} = 0 \quad \text{and}$$

$$(45) \quad (\forall t^1 \overset{\beta}{\approx} t^2) \lim_n \frac{\prod_{a \in t^2} (\rho \cup \sigma_n)(a)}{\prod_{a \in t^1} (\rho \cup \sigma_n)(a)} = \frac{\beta(t^2)}{\beta(t^1)} \in (0, \infty) .$$

The lemma's conclusion follows from (41)–(45) and the definitions of \succ and \approx . \square

Proof B.5 (for Theorem B). Take an assessment (σ, β) and its plausibility relation \succcurlyeq . Assume (σ, β) is consistent and let $\langle (\sigma_n, \beta_n) \rangle_n$ be a sequence of full-support Bayesian assessments which converges to it.

This paragraph shows that there is no pair (s^m, t^m) from \succ in a finite sequence $\langle (s^m, t^m) \rangle_{m=1}^M$ of pairs from \succcurlyeq whenever the sequence satisfies $(\forall a) |\{m | a \in s^m\}| = |\{m | a \in t^m\}|$. Accordingly, let $\langle (s^m, t^m) \rangle_{m=1}^M$ be such a sequence. By the equalities over actions a we have

$$(\forall n) \prod_{m=1}^M \prod_{a \in s^m} (\rho \cup \sigma_n)(a) = \prod_{m=1}^M \prod_{a \in t^m} (\rho \cup \sigma_n)(a) ,$$

which is equivalent to

$$(\forall n) \prod_{m=1}^M \frac{\prod_{a \in t^m} (\rho \cup \sigma_n)(a)}{\prod_{a \in s^m} (\rho \cup \sigma_n)(a)} = 1 ,$$

which obviously yields

$$\lim_n \prod_{m=1}^M \frac{\prod_{a \in t^m} (\rho \cup \sigma_n)(a)}{\prod_{a \in s^m} (\rho \cup \sigma_n)(a)} = 1 .$$

Yet by Lemma B.4 we have

$$\begin{aligned} (\forall s^m \succ t^m) \lim_n \frac{\prod_{a \in t^m} (\rho \cup \sigma_n)(a)}{\prod_{a \in s^m} (\rho \cup \sigma_n)(a)} &= 0 \quad \text{and} \\ (\forall s^m \approx t^m) \lim_n \frac{\prod_{a \in t^m} (\rho \cup \sigma_n)(a)}{\prod_{a \in s^m} (\rho \cup \sigma_n)(a)} &\in (0, \infty) . \end{aligned}$$

The last two sentences contradict if $\langle (s^m, t^m) \rangle_{m=1}^M$ has a pair from \succ . Hence no such pair exists.

Since the previous paragraph has derived condition (b) of Lemma B.3, that lemma provides the existence of $\pi: A \rightarrow \mathbb{Z}$ such that for all s and t ,

$$(46) \quad \begin{aligned} s \succ t &\Rightarrow \sum_{a \in s} \pi(a) > \sum_{a \in t} \pi(a) \quad \text{and} \\ s \approx t &\Rightarrow \sum_{a \in s} \pi(a) = \sum_{a \in t} \pi(a) . \end{aligned}$$

It only remains to show that π is nonpositive-valued. If this were not the case, there would be some a for which $\pi(a) > 0$. Then, since $\{F(h) | h\}$ partitions A by Lemma A.2, there is some t for which $a \in F(t)$. Since $a \notin t$ by the definition of F and since $\pi(a) > 0$ by assumption, we have that $\sum_{a' \in t} \pi(a') < \sum_{a' \in t \cup \{a\}} \pi(a')$, and thus by (46), we have that both $t \overset{\sigma}{\succ} t \cup \{a\}$ and $t \overset{\sigma}{\approx} t \cup \{a\}$ are false. Yet by $a \in F(t)$ and the

definitions of $\overset{\circ}{\approx}$ and $\overset{\circ}{\simeq}$, it must be that either $t \overset{\circ}{\simeq} t \cup \{a\}$ or $t \overset{\circ}{\approx} t \cup \{a\}$ is true. The last two sentences contradict, and hence there cannot be an a for which $\pi(a) > 0$. \square

B.4. ADDITIVE SEPARABILITY

Although this lemma could be proved directly from the definition of a set-tree game, it is quicker to use Theorem A's isomorphism.

Lemma B.6. $(\forall t)(\forall h) |t \cap F(h)| \in \{0, 1\}$. *In other words, a node t can contain no more than one element from each $F(h)$.*

Proof. Consider any $(A, T, H, I, i^c, \rho, u)$. By Theorem A, this set-tree game is isomorphic to a sequence-tree game $(A, \bar{T}, \bar{H}, \bar{I}, \bar{i}^c, \rho, \bar{u})$ with agent recall.

Now suppose there are t and h such that $|t \cap F(h)| > 1$. Then there are distinct a and a' such that $\{a, a'\} \subseteq t \cap F(h)$. First let $\bar{h} = (R_1|_{\mathcal{P}(\bar{T})})^{-1}(h)$, and note that $\{a, a'\} \subseteq \bar{F}(\bar{h})$ because $\bar{F}(\bar{h}) = F(h)$ by Lemma A.3(c). Second let $\bar{t} = (R|_{\bar{T}})^{-1}(t)$, and note that there are distinct m and n such that $\bar{t}_m = a$ and $\bar{t}_n = a'$. By the last two sentences there are distinct m and n such that $\{\bar{t}_m, \bar{t}_n\} \in \bar{F}(\bar{h})$. By Lemma A.4(d) this violates agent recall. \square

Proof B.7 (for Lemma 3.2). V is well-defined because Lemma B.6 shows that $t \cap F(h)$ has no more than one element for any t and any h .

It remains to show that, for any t , t is equal to $\{\dot{t}_h | \dot{t}_h \in A\}$ evaluated at $\dot{t} = V(t)$. Accordingly take any t and note that

$$\begin{aligned} & \{ \dot{t}_h \mid \dot{t}_h \in A \} \text{ evaluated at } \dot{t} = V(t) \\ &= \{ [V(t)]_h \mid [V(t)]_h \in A \} \\ &= \{ a \mid (\exists h) a = [V(t)]_h \} \\ &= \{ a \mid (\exists h) \{a\} = t \cap F(h) \} \\ &= \bigcup_h (t \cap F(h)) \\ &= t, \end{aligned}$$

where the first two equalities follow from manipulation, the third from the definition of $V(t)$, the fourth from Lemma B.6, and the last from Lemma A.2. \square

Proof B.8 (for Corollary 1). Take any consistent assessment. By Theorem B, there exists $\pi:A\rightarrow\mathbb{Z}_-$ to satisfy (8). Since A is partitioned by $\{F(h)|h\}$ by Lemma A.2, we may define $\langle\dot{\pi}_h:\dot{A}_h\rightarrow\mathbb{Z}_-\rangle_h$ at each h by

$$\dot{\pi}_h(\dot{a}_h) = \begin{pmatrix} 0 & \text{if } \dot{a}_h=o \\ \pi(\dot{a}_h) & \text{if } \dot{a}_h\in F(h) \end{pmatrix} .$$

By $\dot{A}_h = \{o\}\cup F(h)$, by the definition of $\dot{\pi}_h$, by $F(h) = \dot{A}_h\cap A$, and by the second half of Lemma 3.2, we have that for all $\dot{t}\in V(T)$

$$\begin{aligned} \Sigma_h\dot{\pi}_h(\dot{t}_h) &= \Sigma_{h|\dot{t}_h=o}\dot{\pi}_h(\dot{t}_h) + \Sigma_{h|\dot{t}_h\in F(h)}\dot{\pi}_h(\dot{t}_h) \\ &= \Sigma_{h|\dot{t}_h=o}0 + \Sigma_{h|\dot{t}_h\in F(h)}\pi(\dot{t}_h) \\ &= \Sigma_{h|\dot{t}_h\in A}\pi(\dot{t}_h) \\ &= \Sigma_{a\in\{\dot{t}_h|\dot{t}_h\in A\}}\pi(a) \\ &= \Sigma_{a\in V^{-1}(\dot{t})}\pi(a) . \end{aligned}$$

By the definition of \succ , by (8), and by the previous sentence we have that for all t^1 and t^2

$$\begin{aligned} t^1 \succ t^2 &\Rightarrow V^{-1}(t^1) \succ V^{-1}(t^2) \\ &\Rightarrow \Sigma_{a\in V^{-1}(t^1)}\pi(a) > \Sigma_{a\in V^{-1}(t^2)}\pi(a) \\ &\Rightarrow \Sigma_h\dot{\pi}_h(t_h^1) > \Sigma_h\dot{\pi}_h(t_h^2) . \end{aligned}$$

Identical reasoning shows $t^1 \approx t^2$ implies $\Sigma_h\dot{\pi}_h(t_h^1) = \Sigma_h\dot{\pi}_h(t_h^2)$. \square

B.5. PROOF OF KREPS-WILSON LEMMA A1

Proof B.9 (for Corollary 2 = KW Lemma A1). The first paragraph of KW's proof shows that a labelled basis must be consistent.

Conversely, suppose b is consistent. By definition this means $\Psi_b \neq \emptyset$, and thus there is a consistent assessment (σ, β) such that

$$(47a) \quad b\cap A = \{a|\sigma(a)>0\} \text{ and}$$

$$(47b) \quad b\cap(T\sim Z) = \{t|\beta(t)>0\} .$$

Let \succ be (π, β) 's plausibility relation. By Theorem B, there exists $\pi:A\rightarrow\mathbb{Z}_-$ such that for all t^1 and t^2

$$(48a) \quad t^1 \succ t^2 \Rightarrow \Sigma_{a\in t^1}\pi(a) > \Sigma_{a\in t^2}\pi(a) \text{ and}$$

$$(48b) \quad t^1 \approx t^2 \Rightarrow \Sigma_{a\in t^1}\pi(a) = \Sigma_{a\in t^2}\pi(a) .$$

Set $K = -\pi$. We will show that K labels b . Since K is nonnegative-valued as required, only (12a–c) remain.

We begin with (12b). Take any a and, by Lemma A.2, let t be such that $a \in F(t)$. First, if $a \in b$, then $\sigma(a) > 0$ by (47a), which implies $t \stackrel{\sigma}{\approx} t \cup \{a\}$ by the definition of $\stackrel{\sigma}{\approx}$, which implies $t \approx t \cup \{a\}$, which implies $\Sigma_{a' \in t} \pi(a') = \Sigma_{a' \in t \cup \{a\}} \pi(a')$ by (48b), which implies $\pi(a) = 0$ by $a \notin t$ by the definition of F , which implies $K(a) = 0$. Conversely, if $a \notin b$, then $\sigma(a) = 0$ by (47a), which implies $t \stackrel{\sigma}{\succ} t \cup \{a\}$ by the definition of $\stackrel{\sigma}{\succ}$, which implies $t \succ t \cup \{a\}$, which implies $\Sigma_{a' \in t} \pi(a') > \Sigma_{a' \in t \cup \{a\}} \pi(a')$ by (48a), which implies $\pi(a) < 0$, which implies $K(a) > 0$, which implies $K(a) \neq 0$.

Next we turn to (12a). Since $\sigma|_{F(h)}$ is a probability distribution, there is some $a \in F(h)$ such that $\sigma(a) > 0$. This implies $a \in b$ by (47a), and hence $K(a) = 0$ by (12b), which has already been proved by the previous paragraph.

Finally, turn to (12c). Take any h and any $t \in h$. First, suppose $t \in b$. Then $\beta(t) > 0$ by (47b). Now consider any other $t' \in h$. By the definitions of $\stackrel{\beta}{\succ}$ and $\stackrel{\beta}{\approx}$, either $t \stackrel{\beta}{\succ} t'$ or $t \stackrel{\beta}{\approx} t'$, and thus in either event we have $t \succcurlyeq t'$. Hence by (48), $\Sigma_{a \in t} \pi(a) \geq \Sigma_{a \in t'} \pi(a)$, which implies $\Sigma_{a \in t} K(a) \leq \Sigma_{a \in t'} K(a)$, which implies $J_K(t) \leq J_K(t')$. Since this holds for any other $t' \in h$, we have that $t \in \operatorname{argmin}\{J_K(t') | t' \in h\}$. Conversely, suppose $t \notin b$. Then $\beta(t) = 0$ by (47b). Since $\beta|_h$ is a probability distribution, there is some t^* such that $\beta(t^*) > 0$. Thus $t^* \stackrel{\beta}{\succ} t$ by the definition of $\stackrel{\beta}{\succ}$, which implies $t^* \succ t$, which implies $\Sigma_{a \in t^*} \pi(a) > \Sigma_{a \in t} \pi(a)$ by (48a), which implies $\Sigma_{a \in t^*} K(a) < \Sigma_{a \in t} K(a)$, which implies $J_K(t^*) < J_K(t)$, which implies $t \notin \operatorname{argmin}\{J_K(t') | t' \in h\}$. \square

REFERENCES

- ANDERLINI, L., D. GERARDI, AND R. LAGUNOFF (2008): "A Super Folk Theorem for Dynastic Repeated Games," *Economic Theory*, 37, 357–394.
- BATTIGALLI, P. (1996): "Strategic Independence and Perfect Bayesian Equilibrium," *Journal of Economic Theory*, 70, 201–234.
- BLACKORBY, C., D. PRIMONT, AND R. R. RUSSELL (1978): *Duality, Separability, and Functional Structure: Theory and Economic Applications*. North-Holland.
- BLUME, L., A. BRANDENBURGER, AND E. DEKEL (1991a): "Lexicographic Probabilities and Choice Under Uncertainty," *Econometrica*, 59, 61–79.
- (1991b): "Lexicographic Probabilities and Equilibrium Refinements," *Econometrica*, 59, 81–98.
- BLUME, L. E., AND W. R. ZAME (1994): "The Algebraic Geometry of Perfect and Sequential Equilibrium," *Econometrica*, 62, 783–794.

- DEBREU, G. (1960): "Topological Methods in Cardinal Utility Theory," in *Mathematical Methods in the Social Sciences, 1959*, ed. by K. J. Arrow, S. Karlin, and P. Suppes, pp. 16–26. Stanford University Press, reprinted in Debreu (1983).
- (1983): *Mathematical Economics: Twenty Papers of Gerard Debreu*. Cambridge University Press.
- FARKAS, J. (1902): "Über die Theorie der einfachen Ungleichungen," *J. Math.*, 124, 1–24.
- FUDENBERG, D., AND D. K. LEVINE (2006): "Superstition and Rational Learning," *American Economic Review*, 96, 630–651.
- GORMAN, W. M. (1968): "The Structure of Utility Functions," *Review of Economic Studies*, 35, 376–390.
- GOVINDAN, S., AND R. WILSON (2001): "Direct Proofs of Generic Finiteness of Nash Equilibrium Outcomes," *Econometrica*, 69, 765–769.
- HALPERN, J. Y. (2010): "Lexicographic probability, conditional probability, and nonstandard probability," *Games and Economic Behavior*, 68, 155–179.
- HAMMOND, P. J. (1994): "Elementary Non-Archimedean Representations of Probability for Decision Theory and Games," in *Patrick Suppes: Scientific Philosopher, Volume 1*, ed. by P. Humphreys, pp. 25–61. Kluwer.
- KOBAYASHI, H. (2007): "Folk Theorems for Infinitely Repeated Games Played by Organizations with Short-Lived Members," *International Economic Review*, 48, 407–549.
- KOHLBERG, E., AND P. J. RENY (1997): "Independence on Relative Probability Spaces and Consistent Assessments in Game Trees," *Journal of Economic Theory*, 75, 280–313.
- KRAFT, C. H., J. W. PRATT, AND A. SEIDENBERG (1959): "Intuitive Probability on Finite Sets," *Annals of Mathematical Statistics*, 30, 408–419.
- KRANTZ, D. H., R. D. LUCE, P. SUPPES, AND A. TVERSKY (1971): *Foundations of Measurement, Volume I: Additive and Polynomial Representations*. Academic Press.
- KREPS, D. M., AND G. RAMEY (1987): "Structural Consistency, Consistency, and Sequential Rationality," *Econometrica*, 55, 1331–1348.
- KREPS, D. M., AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.
- KUHN, H. W. (1953): "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games*, ed. by H. W. Kuhn, and A. W. Tucker, pp. 193–216. Princeton University Press, reprinted in Kuhn (1997).
- (1997): (editor) *Classics in Game Theory*. Princeton University Press.
- MAS-COLELL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*. Oxford University Press.
- MCLENNAN, A. (1989): "Consistent Conditional Systems in Noncooperative Game Theory," *International Journal of Game Theory*, 18, 141–174.
- MYERSON, R. B. (1986): "Multistage Games with Communication," *Econometrica*, 54, 323–358.
- NARENS, L. (1985): *Abstract Measurement Theory*. MIT Press.
- OSBORNE, M. J. (2008): "Strategic and Extensive Form Games," in *The New Palgrave Dictionary of Economics, Second Edition*, ed. by S. N. Durlauf, and L. E. Blume. Palgrave Macmillan.

- OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT Press.
- PEREA, A. (2001): *Rationality in Extensive Form Games*. Kluwer.
- PEREA Y MONSUWÉ, A., M. JANSEN, AND H. PETERS (1997): “Characterization of Consistent Assessments in Extensive Form Games,” *Games and Economic Behavior*, 21, 238–252.
- PICCIONE, M., AND A. RUBINSTEIN (1997): “On the Interpretation of Decisions Problems with Imperfect Recall,” *Games and Economics Behavior*, 20, 3–24.
- SCOTT, D. (1964): “Measurement Structures and Linear Inequalities,” *Journal of Mathematical Psychology*, 1, 233–247.
- STREUFERT, P. A. (2006): “A Comment on ‘Sequential Equilibria’,” University of Western Ontario, Economics Department Research Report 2006-03, 10 pages.
- VON STENGEL, B. (1996): “Efficient Computation of Behavior Strategies,” *Games and Economic Behaviour*, 14, 220–246.