

P. GALLIANI
A. L. MANN

Lottery Semantics: A Compositional Semantics for Probabilistic First-Order Logic with Imperfect Information

Abstract. We present a compositional semantics for first-order logic with imperfect information that is equivalent to Sevenster and Sandu's equilibrium semantics (under which the truth value of a sentence in a finite model is equal to the minimax value of its semantic game). Our semantics is a generalization of an earlier semantics developed by the first author that was based on behavioral strategies, rather than mixed strategies.

Keywords: Logic with imperfect information, independence-friendly logic, dependence-friendly logic, dependence logic, equilibrium semantics, compositional semantics.

Introduction

Game-theoretic semantics is an approach to first-order logic that defines truth and satisfaction in terms of (semantic) games. Although the fundamental intuition that quantifiers can be interpreted as moves in a game is present already in Peirce's second Cambridge Conferences lecture [18], as well as Henkin's seminal paper on branching quantifiers [7], game-theoretic semantics was first popularized by Hintikka [8, 9] (see also [11, 13]).

In brief, the semantic game associated with a first-order sentence is a contest between two opponents. One tries to verify the sentence by choosing the values of existentially quantified variables, while the other tries to falsify it by picking the values of universally quantified variables. Disjunctions prompt the existential player to choose a disjunct; conjunctions prompt the universal player to pick a conjunct. Negation tells the players to switch roles. A first-order sentence is true (false) in a model if and only if the existential (universal) player has a winning strategy.

In order to define the semantic game for an open¹ first-order formula, one must specify the values of its free variables. Usually, this is done using an assignment. If the open formula in question is a subformula of some

Presented by **Name of Editor**; *Received* December 1, 2005

¹A first-order formula is said to be *open* if it has free variables, that is, if it is not a sentence.

first-order sentence, then we can think of the assignment as encoding the previous moves of the players in the semantic game for the sentence.

In the semantic game for any first-order formula, the players take turns making their moves, and at each decision point the active player is aware of every move leading up to the current position. Such games can thus be modeled as extensive games with perfect information.

First-order logic with imperfect information is an extension of first-order logic obtained by considering semantic games with imperfect information. In a game with imperfect information, the active player may not be aware of every move leading up to the current position. To specify such games, we must extend the syntax of first-order logic to be able to indicate what information is available to the active player. We briefly describe three approaches found in the literature.

Independence-friendly (IF) logic, introduced by Hintikka and Sandu [12], adds a slash set to each quantifier that indicates which variables the active player is not allowed to access when choosing the value of the quantified variable. For example, in the independence-friendly sentence

$$\forall x(\exists y/\{x\})Rxy,$$

the existential player must choose the value of y without knowing the value of x . One drawback of IF logic is that by specifying which variables a player is *not* allowed to see, the set of variables the player *is* allowed to see depends not only on the quantifier itself, but also on which variables have been assigned values.

An alternative approach, called dependence-friendly (DF) logic, specifies which variables the active player is allowed to access. Traditionally, the set of variables whose values the active player can see is indicated using a backslash instead of a forward slash. In an attempt to lighten our notation, we will instead place such sets in a superscript above the relevant quantifier. For example, the dependence-friendly sentence

$$\forall x^\emptyset \exists y^\emptyset Rxy$$

has the same semantic game as the IF sentence above because the values of x and y are chosen independently.

Väänänen goes a step further by introducing new atomic formulas of the form

$$=(t_1, \dots, t_n)$$

whose intuitive meaning is that the value of the term t_n depends only on the values of the terms t_1, \dots, t_{n-1} . The atomic formula $=(t)$ asserts that

the value of t is constant [21, p. 17]. Thus, when playing the semantic game for the dependence logic formula

$$\forall x \exists y (= (y) \wedge Rxy),$$

the existential player knows the value of x when choosing the value of y , but if the game is repeated she must choose the same value for y as before, regardless of the new value of x .

All three variants have the same expressive power as existential second-order logic, a result first proved independently by Enderton [3] and Walkoe [23] for first-order logic with branching quantifiers. In the present paper it will be convenient to work with dependence-friendly logic, but all of our results can be adapted to the other two frameworks.

1. Syntax

DEFINITION 1.1. For a fixed vocabulary L , the set of L -terms and the set of atomic formulas with vocabulary L are defined as for first-order logic. The set DF_L of dependence-friendly formulas with vocabulary L is generated by the grammar

$$\chi \mid \sim\phi \mid (\phi \vee \psi) \mid \exists x^V \phi$$

where χ ranges over all atomic formulas with vocabulary L , x ranges over the variables in a countably infinite set $\{x_1, x_2, \dots\}$, and V ranges over finite subsets of $\{x_1, x_2, \dots\}$. We adopt the standard abbreviations of $(\varphi \wedge \psi)$ for $\sim(\sim\varphi \vee \sim\psi)$ and $\forall x^V \phi$ for $\sim\exists x^V \sim\phi$.

The set V in the formula $\exists x^V \phi$ is called a *dependence set* because it specifies upon which variables the value of x depends. In principle, we could attach dependence sets to disjunctions (and conjunctions), but we prefer not to in order to simplify the presentation. Adding dependence sets to disjunctions would not increase the expressive power of the logic because we can simply take $\phi \vee^W \psi$ to be an abbreviation for

$$\exists z^W [(z = 0 \wedge \phi) \vee (z = 1 \wedge \psi)],$$

where 0 and 1 are constant symbols, and z is a fresh variable that does not occur in ϕ , ψ , or W .

Free variables are defined as usual, except that in a formula of the form $\exists x^V \phi$ the variables in V are considered to be free.

DEFINITION 1.2. The set $\text{Free}(\phi)$ of *free variables* of a DF formula ϕ is defined recursively. If ϕ is atomic every variable is free. In addition,

- $\text{Free}(\sim\phi) = \text{Free}(\phi)$,
- $\text{Free}(\phi \vee \psi) = \text{Free}(\phi) \cup \text{Free}(\psi)$,
- $\text{Free}(\exists x^V \phi) = (\text{Free}(\phi) \setminus \{x\}) \cup V$.

Notice that the variable x can be free in a formula of the form $\exists x^V \phi$ or $\forall x^V \phi$.

2. Game-theoretic semantics

For the sake of comparison, we first define the semantic game for first-order formulas. One can imagine that after every move of the semantic game, a new subgame begins. Thus we need a way to encode the state of the game in such a way that the subgame starts in the correct position. The state of the semantic game for a first-order sentence consists of the current subformula, the values of the variables, and which player is the verifier.

DEFINITION 2.1. An *assignment* is a function that assigns values to variables. If \mathbf{M} is a structure with universe M , and V is a set of variables, then a function $s: V \rightarrow M$ is called an *\mathbf{M} -valued assignment*. If $m \in M$, the assignment $s[x/m]$ is defined by

$$s[x/m](y) = \begin{cases} m & \text{if } x = y, \\ s(y) & \text{if } x \neq y. \end{cases}$$

If $V' \subseteq V$, then

$$s|_{V'} = \left\{ \langle x, s(x) \rangle : x \in V' \right\}$$

is called the *restriction of s to V'* .

DEFINITION 2.2. Let ϕ be a first-order formula, \mathbf{M} a suitable structure, and s an \mathbf{M} -valued assignment whose domain contains $\text{Free}(\phi)$. The *semantic game* $G(\mathbf{M}, s, \phi)$ is defined as follows.

- A *position* of the game is a triple $\langle \psi, s', \alpha \rangle$, where ψ is a particular instance² of a subformula of ϕ , s' is an assignment whose domain contains $\text{Free}(\psi)$, and $\alpha \in \{\exists, \forall\}$. We invite Eloise to play the role of the existential player and Abelard to play the role of the universal player. When ψ is an atomic formula, $\langle \psi, s', \alpha \rangle$ is called a *terminal position*.
- The set of all positions of the game is denoted by P . Let P_α denote the set of non-terminal positions of the form $\langle \psi, s', \alpha \rangle$.

²For example, if ϕ is $\psi \vee \psi$, we distinguish between the left and right disjuncts.

- A *history* or *play* is a finite sequence of positions p_0, \dots, p_n such that for all $0 \leq i < n$:
 - $p_0 = \langle \phi, s, \exists \rangle$;
 - if $p_i = \langle \sim\psi, s', \alpha \rangle$ then $p_{i+1} = \langle \psi, s', \bar{\alpha} \rangle$, where $\bar{\exists} = \forall$ and $\bar{\forall} = \exists$;
 - if $p_i = \langle \psi \vee \chi, s', \alpha \rangle$ then $p_{i+1} = \langle \psi, s', \alpha \rangle$ or $p_{i+1} = \langle \chi, s', \alpha \rangle$;
 - if $p_i = \langle \exists x^V \psi, s', \alpha \rangle$ then $p_{i+1} = \langle \psi, s'[x/m], \alpha \rangle$ for some $m \in M$;
 - if p_i is a terminal position, then $i = n$.
- A *terminal history* or *complete play* is a history p_0, \dots, p_n such that $p_n = \langle \chi, s', \alpha \rangle$ is a terminal position. The winner of a terminal history is
 - player α if $\mathbf{M}, s' \models \chi$,
 - player $\bar{\alpha}$ if $\mathbf{M}, s' \not\models \chi$.

When the assignment s is empty we simply write $G(\mathbf{M}, \phi)$.

A *strategy* for player α in $G(\mathbf{M}, s, \phi)$ is a function $\sigma: P_\alpha \rightarrow P$ such that if p_0, \dots, p_n is a history, and $p_n \in P_\alpha$, then $p_0, \dots, p_n, \sigma(p_n)$ is also a history. Player α is said to *follow* a strategy σ in a history p_0, \dots, p_n if for all $0 \leq i < n$ such that $p_i \in P_\alpha$ we have $\sigma(p_i) = p_{i+1}$. A strategy is *winning* if its owner wins every complete play in which he or she follows it. The Gale-Stewart theorem [4] implies that for any first-order semantic game, either Eloise or Abelard must have a winning strategy.

DEFINITION 2.3. An assignment s *satisfies* a first-order formula ϕ in a model \mathbf{M} , written $\mathbf{M} \models_s \phi$, if Eloise has a winning strategy for the semantic game $G(\mathbf{M}, s, \phi)$. A first-order sentence ϕ is *true* in \mathbf{M} , written $\mathbf{M} \models \phi$, if Eloise has a winning strategy for $G(\mathbf{M}, \phi)$.

Semantic games for DF formulas are similar to those for first-order formulas except that the players do not always have access to the entire assignment when making their moves. Thus, in order to encode the state of the game during play, it is no longer sufficient to simply record the values of the variables. We must also record what the players know about the values of the variables.

DEFINITION 2.4. A *team* is a set of assignments with the same domain. A team of \mathbf{M} -valued assignments is called an *\mathbf{M} -valued team*.

DEFINITION 2.5. Let ϕ be a DF formula, \mathbf{M} a suitable structure, and X a team whose domain contains $\text{Free}(\phi)$. The *semantic game* $G(\mathbf{M}, X, \phi)$ is defined as above, except:

- The initial position of a history p_0, \dots, p_n may be any member of

$$\{ \langle \phi, s, \exists \rangle : s \in X \}.$$

- Two positions $\langle \exists x^V \psi, s, \alpha \rangle$ and $\langle \exists x^V \psi, s', \alpha \rangle$ are *indistinguishable* if $s|_V = s'|_V$ (in which case we write $s =_V s'$).

When the team X includes only the empty assignment, we simply write $G(\mathbf{M}, \phi)$.

A team encodes the information that the players have about the current assignment at the beginning of the game. More precisely, a team X represents the knowledge that the current assignment belongs to X . The semantic game for a DF formula is unusual in that it does not have a unique starting position. It may help the reader's intuition to imagine that, at the beginning of the game, the initial assignment is chosen from X by a disinterested third party (Nature), after which Eloise makes her first move.

A strategy σ for player α in the game $G(\mathbf{M}, X, \phi)$ must be *uniform* in the sense that if two positions $p = \langle \exists x^V \psi, s, \alpha \rangle$ and $p' = \langle \exists x^V \psi, s', \alpha \rangle$ are indistinguishable, then $\sigma(p) = \sigma(p')$. In general, the lack of perfect information about the current assignment prevents the players from following strategies that would otherwise be available, but it does not prevent them from performing any particular action. For example, in the semantic game for

$$\forall x^{\emptyset} \exists y^{\emptyset} Rxy$$

played in the structure $\mathbf{2} = \{0, 1\}$, each player has two possible strategies: $x := 0$ or $x := 1$ for Abelard, and $y := 0$ or $y := 1$ for Eloise. In contrast, in the semantic game for the first-order sentence

$$\forall x \exists y Rxy$$

Eloise may follow the additional strategies $y := x$ and $y := 1 - x$.

The fact that Eloise and Abelard may be prevented from following certain strategies means that it is possible for neither of them to have a winning strategy. For example, none of the uniform strategies for $\forall x^{\emptyset} \exists y^{\emptyset} x = y$ in a structure with at least two elements is winning. Hence the sentence is neither true nor false.

DEFINITION 2.6. A team X *satisfies* a DF formula ϕ in a model \mathbf{M} , written $\mathbf{M} \models_X^+ \phi$, if Eloise has a winning strategy for $G(\mathbf{M}, X, \phi)$, in which case we also say that X is a *winning team* for ϕ in \mathbf{M} . A team X *dissatisfies*

ϕ , written $\mathbf{M} \models_X^- \phi$, if Abelard has a winning strategy for $G(\mathbf{M}, X, \phi)$, in which case we say that X is a *losing team* for ϕ in \mathbf{M} .

A DF sentence ϕ is *true* in \mathbf{M} , written $\mathbf{M} \models^+ \phi$, if Eloise has a winning strategy for $G(\mathbf{M}, \phi)$, and it is *false* in \mathbf{M} , written $\mathbf{M} \models^- \phi$, if Abelard has a winning strategy.

3. Trump semantics

Tarski's semantics for first-order logic is a recursive procedure for analyzing the semantic game of a first-order sentence in terms of its subgames. Originally, the game-theoretic semantics for IF logic was only defined for sentences [10, 12]. Later Hodges extended Hintikka and Sandu's semantics to open IF formulas, and showed how to analyze the semantic game of an IF sentence in terms of its subgames. Hodges called a winning team a *trump*, and a losing team a *cotrump*; hence the name *trump semantics* [14, 15].

In order to pass from game to subgame, we need a way to update the information the players have about the current assignment. When Eloise chooses the value of an existentially quantified variable, she can calculate the effect of her choice on all the assignments she considers possible. Given a team X and a set of variables V , let

$$X|_V = \{ s|_V : s \in X \}.$$

If $F: X|_V \rightarrow M$ and $s \in X$, we abuse notation by writing $F(s) = F(s|_V)$. Define the *supplement team*

$$X[x/F] = \{ s[x/F(s)] : s \in X \}.$$

Even when he is completely ignorant of the current assignment, Abelard is free to choose any element of the universe as the value of a universally quantified variable. Hence Eloise cannot assume anything about the element he picked. Define the *duplicate team*

$$X[x/M] = \{ s[x/m] : s \in X, m \in M \}.$$

THEOREM 3.1 (Hodges [14, Theorem 7.5]). *Let \mathbf{M} be a structure, and let X be a team. First, the clauses for satisfaction:*

- If ϕ is atomic, then $\mathbf{M} \models_X^+ \phi$ if and only if for all $s \in X$, $\mathbf{M} \models_s \phi$.
- $\mathbf{M} \models_X^+ \sim\phi$ if and only if $\mathbf{M} \models_X^- \phi$.

- $\mathbf{M} \models_X^+ \phi \vee \psi$ if and only if there exists a cover $X = Y \cup Z$ such that

$$\mathbf{M} \models_Y^+ \phi \quad \text{and} \quad \mathbf{M} \models_Z^+ \psi.$$

- $\mathbf{M} \models_X^+ \exists x^V \phi$ if and only if there is a function $F: X|_V \rightarrow M$ such that

$$\mathbf{M} \models_{X[x/F]}^+ \phi.$$

Now the clauses for dissatisfaction:

- If ϕ is atomic, then $\mathbf{M} \models_X^- \phi$ if and only if for all $s \in X$, $\mathbf{M} \not\models_s \phi$.
- $\mathbf{M} \models_X^- \sim \phi$ if and only if $\mathbf{M} \models_X^+ \phi$.
- $\mathbf{M} \models_X^- \phi \vee \psi$ if and only if $\mathbf{M} \models_X^- \phi$ and $\mathbf{M} \models_X^- \psi$.
- $\mathbf{M} \models_X^- \exists x^V \phi$ if and only if $\mathbf{M} \models_{X[x/M]}^- \phi$.

4. Equilibrium semantics

A DF sentence can be true, false, or neither. Thus, considering semantic games with imperfect information introduces a third truth value: *undetermined*.


So far, we have treated all undetermined DF sentences the same. But if neither player has a winning strategy, which strategies should the players follow? Intuitively, Eloise should follow a strategy that allows her to win as often as possible, regardless of the fact that she cannot win every play. If Eloise always follows the same non-winning strategy σ , there must be a strategy τ for Abelard such that he wins the play determined by σ and τ . But since τ is non-winning there must be a strategy σ' for Eloise such that she wins the play determined by σ' and τ . For example, in the game

$$G(\mathbf{2}, \exists x^\emptyset \forall y^\emptyset x = y),$$

the strategy $x := 0$ wins against the strategy $y := 0$, but loses against $y := 1$. However, $x := 1$ beats $y := 1$, and so on (see Figure 1).

In order to stop going around in circles, we must allow the players to randomize their strategies.

DEFINITION 4.1. Let ϕ be a DF sentence, and let \mathbf{M} be a suitable structure. Let S_\exists denote the set of Eloise's (pure) strategies for the semantic game $G(\mathbf{M}, \phi)$, and let S_\forall denote the set of Abelard's (pure) strategies. We say that $G(\mathbf{M}, \phi)$ is *finite* if S_\exists and S_\forall are both finite.


 Figure 1. The semantic game for $\exists x^0 \forall y^0 x = y$

Since a pair of pure strategies $\langle \sigma, \tau \rangle \in S_{\exists} \times S_{\forall}$ determines a complete play of the game, we can define the *utility for player α* by

$$u_{\alpha}(\sigma, \tau) = \begin{cases} 1 & \text{if } \alpha \text{ wins,} \\ 0 & \text{if } \bar{\alpha} \text{ wins.} \end{cases}$$

A *mixed strategy* for player α is a probability distribution over S_{α} , the set of which is denoted $\Delta(S_{\alpha})$. Given a pair of mixed strategies $\langle \mu, \nu \rangle \in \Delta(S_{\exists}) \times \Delta(S_{\forall})$, the *expected utility for player α* is

$$U_{\alpha}(\mu, \nu) = \int \int u_{\alpha} d\mu d\nu.$$

When S_{\exists} and S_{\forall} are both finite, we have

$$U_{\alpha}(\mu, \nu) = \sum_{\sigma \in S_{\exists}} \sum_{\tau \in S_{\forall}} \mu(\sigma) \nu(\tau) u_{\alpha}(\sigma, \tau).$$

Since $u_{\bar{\alpha}} = 1 - u_{\alpha}$ and $U_{\bar{\alpha}} = 1 - U_{\alpha}$, we will focus on Eloise's utility function $u = u_{\exists}$ and expected utility function $U = U_{\exists}$.

EXAMPLE 4.2. Let us revisit the game $G(\mathbf{2}, \exists x^0 \forall y^0 x = y)$. If Eloise follows the strategy $x := 0$ with probability p , and Abelard follows the strategy $y := 0$ with probability q , then Eloise's expected utility is

$$U(\mu_p, \nu_q) = pq + (1 - p)(1 - q).$$

Hence, if $p = 1/3$ and $q = 3/4$, Eloise's expected utility is $5/12$.

Notice that if we fix Abelard's mixed strategy of playing $y := 0$ with probability $3/4$, Eloise's expected utility is

$$U(\mu_p, \nu_{3/4}) = \frac{3}{4}p + \frac{1}{4}(1-p) = \frac{1}{2}p + \frac{1}{4}.$$

Thus Eloise maximizes her expected utility by always playing $x := 0$.

By allowing the players to use mixed strategies we can assign intermediate truth values to undetermined DF sentences.

DEFINITION 4.3. A pair $\langle \mu^*, \nu^* \rangle \in \Delta(S_{\exists}) \times \Delta(S_{\forall})$ of mixed strategies is a *Nash equilibrium* if

- for all $\mu \in \Delta(S_{\exists})$ we have $U(\mu, \nu^*) \leq U(\mu^*, \nu^*)$,
- for all $\nu \in \Delta(S_{\forall})$ we have $U(\mu^*, \nu^*) \leq U(\mu^*, \nu)$.

PROPOSITION 4.4. *If $\langle \mu, \nu \rangle$ and $\langle \mu', \nu' \rangle$ are both Nash equilibria, then*

$$U(\mu, \nu) = U(\mu', \nu').$$

PROOF. $U(\mu, \nu) \leq U(\mu, \nu') \leq U(\mu', \nu') \leq U(\mu', \nu) \leq U(\mu, \nu)$. □

Proposition 4.4 shows that any two Nash equilibria for a semantic game will yield the same expected utilities for the players. Nash equilibria for semantic games do not always exist, however. Consider the semantic game for the sentence

$$\forall x^0 \exists y^0 x \leq y$$

played on the natural numbers. No pair of mixed strategies $\langle \mu, \nu \rangle \in \Delta(S_{\exists}) \times \Delta(S_{\forall})$ is a Nash equilibrium because each player can improve her or his expected utility by choosing yet greater numbers with higher probability.

Fortunately for us, von Neumann's minimax theorem states that Nash equilibria³ exist for every two-player, zero-sum game in which each player has a finite number of pure strategies [22]. Thus, when \mathbf{M} is finite, $G(\mathbf{M}, \phi)$ has a Nash equilibrium. Furthermore, since any two Nash equilibria yield the same expected utilities, we can define the *value* of the game to be

$$\text{Val}(\mathbf{M}, \phi) = U(\mu^*, \nu^*),$$

where $\langle \mu^*, \nu^* \rangle$ is any Nash equilibrium for $G(\mathbf{M}, \phi)$. In other words, the truth value of a DF sentence ϕ in a finite model \mathbf{M} is the probability that Eloise wins a play of the semantic game $G(\mathbf{M}, \phi)$, assuming the players follow (mixed) strategies that are in equilibrium.

³We adopt the modern terminology even though the minimax theorem precedes Nash's work by 31 years.

DEFINITION 4.5. Let ϕ be a DF sentence and \mathbf{M} a suitable finite structure. We write $\mathbf{M} \models^\varepsilon \phi$ if and only if $\text{Val}(\mathbf{M}, \phi) = \varepsilon$.

EXAMPLE 4.6. Let $\mathbf{n} = \{0, \dots, n-1\}$ be a structure with n elements. Then

$$\mathbf{n} \models^{1/n} \exists x^\emptyset \forall y^\emptyset x = y.$$

Consider the uniform mixed strategies

$$\mu^* = \sum_{i \in \mathbf{n}} \frac{1}{n} (x := i) \quad \text{and} \quad \nu^* = \sum_{j \in \mathbf{n}} \frac{1}{n} (y := j)$$

for Eloise and Abelard, respectively, and observe that $U(\mu^*, \nu^*) = 1/n$. To show $\langle \mu^*, \nu^* \rangle$ is a Nash equilibrium, note that if Eloise follows any other mixed strategy μ ,

$$\begin{aligned} U(\mu, \nu^*) &= \sum_{i \in \mathbf{n}} \sum_{j \in \mathbf{n}} \mu(x := i) \nu^*(y := j) u(x := i, y := j) \\ &= \frac{1}{n} \left[\sum_{i \in \mathbf{n}} \left(\mu(x := i) \sum_{j \in \mathbf{n}} u(x := i, y := j) \right) \right] \\ &= \frac{1}{n} \left(\sum_{i \in \mathbf{n}} \mu(x := i) \right) \\ &= 1/n. \end{aligned}$$

Furthermore, if Abelard switches to any other mixed strategy ν , then

$$U(\mu^*, \nu) = 1/n$$

by a similar calculation. Thus neither player has an incentive to deviate from $\langle \mu^*, \nu^* \rangle$.

The previous example is due to Miklos Ajtai, who suggested using the minimax theorem to assign intermediate truth values to sentences with branching quantifiers [1, p. 16]. Later, Sevenster [19] followed by Sevenster and Sandu [20] implemented Ajtai's suggestion for IF sentences and coined the name *equilibrium semantics*. Definition 4.5 is simply an adaptation of equilibrium semantics to DF logic.

Independently of Sevenster and Sandu, the first author developed a probabilistic semantics for DF logic based on behavioral strategies,⁴ and showed

0, 3	1, 3	2, 3	3, 3
0, 2	1, 2	2, 2	3, 2
0, 1	1, 1	2, 1	3, 1
0, 0	1, 0	2, 0	3, 0

Figure 2. A Nash equilibrium for $G(\mathbf{4}, \exists x^\theta \forall y^\theta x = y)$

how to analyze his semantics compositionally [5, 6]. The goal of the present paper is to do the same for equilibrium semantics.

We can visualize the Nash equilibrium in Example 4.6 as a rectangular partition of the unit square (see Figure 2). The regions where Eloise wins are shaded. Moreover, we can represent any pair of mixed strategies using a similar diagram.

DEFINITION 4.7. Let ϕ be a DF formula, and let \mathbf{M} be a suitable finite structure. Then each player has only finitely many pure strategies for the semantic game $G(\mathbf{M}, \phi)$. A *representation of a mixed strategy* $\mu \in \Delta(S_\alpha)$ is a partition of the unit interval $[0, 1]$ into subintervals A_1, \dots, A_m such that, for each pure strategy σ_i in the support of μ , the length of A_i equals $\mu(\sigma_i)$.

A *representation of a pair of mixed strategies* $\langle \mu, \nu \rangle \in \Delta(S_\exists) \times \Delta(S_\forall)$ is

⁴A *behavioral strategy* tells a player to choose which action to take based on a different probability distribution at each decision point. In contrast, a mixed strategy tells a player to choose which pure strategy to follow according to a probability distribution over the set of all pure strategies. Every behavioral strategy corresponds to a mixed strategy, but not vice versa. Kuhn's theorem states that in a game with perfect recall every mixed strategy corresponds to a behavioral strategy [16].

Not all semantic games have perfect recall, so Kuhn's Theorem does not apply. For example, in the game $G(\mathbf{2}, \exists x^\theta \exists y^\theta x = y)$, the mixed strategy

$$\frac{1}{2}(x := 0, y := 0) + \frac{1}{2}(x := 1, y := 1)$$

does not correspond to any behavioral strategy.

A *behavioral equilibrium* is a Nash equilibrium $\langle \mu^*, \nu^* \rangle$ where both μ^* and ν^* are behavioral strategies. Although behavioral strategies are easier to analyze than mixed strategies, it must be stressed that the minimax theorem does not guarantee the existence of behavioral equilibria.




Figure 3. A representation of a pair of mixed strategies

a partition

$$\{ A_i \times B_j : A_i \in \mathcal{A}, B_j \in \mathcal{B} \}$$

of the unit square $[0, 1]^2$ such that $\mathcal{A} = \{A_1, \dots, A_m\}$ is a representation of μ and $\mathcal{B} = \{B_1, \dots, B_n\}$ is a representation of ν (see Figure 3). The area of the shaded region equals Eloise’s expected utility $U(\mu, \nu)$. Thus, when $\langle \mu, \nu \rangle$ is a Nash equilibrium, the shaded area corresponds to the value of the game.

We will use such representations in the next section to define a compositional semantics for DF logic that is equivalent to equilibrium semantics.

5. Lottery games

The first challenge Hodges faced when trying to develop a compositional semantics for IF formulas was that Hintikka and Sandu only defined their game-theoretic semantics for IF sentences. We now find ourselves in a similar position. To calculate the truth value of a DF sentence (in a finite model) in terms of the truth values of its subformulas, we must extend equilibrium semantics to open formulas.⁵ In other words, we need a way to encode the state of the semantic game after each move.

⁵Mann, Sandu, and Sevenster [17] define equilibrium semantics for open IF formulas, but their semantics is not fully compositional.

At this point, we confront a seemingly insurmountable obstacle. When a player follows a mixed strategy, he or she picks which pure strategy to follow at the beginning of the game, and then executes his or her chosen strategy without further reflexion. In a sense, each player only moves once, and both players move simultaneously. We call the version of the semantic game which ends immediately after both players have chosen their strategies a *strategic DF game*. It is difficult to decompose strategic DF games into their subgames because such games do not have subgames!

In order to analyze strategic DF games, we simulate them using a new kind of game defined below.

DEFINITION 5.1. A *grid* is a pair $\langle \mathcal{A}, \mathcal{B} \rangle$ of finite partitions of the unit interval such that every $A \in \mathcal{A}$ and $B \in \mathcal{B}$ is measurable. A partial function H on the unit square *respects* the grid $\langle \mathcal{A}, \mathcal{B} \rangle$ if it is constant on every rectangle defined by the grid:

$$a, a' \in A \in \mathcal{A} \text{ and } b, b' \in B \in \mathcal{B} \quad \text{imply} \quad H(a, b) = H(a', b').$$

We allow ourselves to write $H(a, b) = H(a', b')$ when both are undefined.

DEFINITION 5.2. A *strategy guide* for a structure \mathbf{M} that assigns values to the variables in V is a partial function $H: [0, 1]^2 \rightarrow M^V$ on the unit square that respects a grid. For each fixed $a \in [0, 1]$ we let

$$H(a, *) = \{ H(a, b) : b \in [0, 1] \text{ and } H(a, b) \text{ is defined} \},$$

and for each fixed $b \in [0, 1]$,

$$H(*, b) = \{ H(a, b) : a \in [0, 1] \text{ and } H(a, b) \text{ is defined} \}.$$

Note that $H(a, b)$ is an assignment, while $H(a, *)$ and $H(*, b)$ are teams. We call $H(a, *)$ a *vertical cross-section* and $H(*, b)$ a *horizontal cross-section*.

Strategy guides will play the role in our compositional semantics that teams play for trump semantics.

DEFINITION 5.3. Let ϕ be a DF formula, \mathbf{M} a suitable finite structure, and H a strategy guide for \mathbf{M} that assigns values to the free variables of ϕ (and possibly other variables). The *lottery game* $\underline{G}(\mathbf{M}, H, \phi)$ is defined as follows.

- The set of positions is defined as in Definition 2.2.
- A *history* or *play* is a finite sequence $\langle a, b \rangle, p_0, \dots, p_n$ such that:
 - $\langle a, b \rangle$ is a point in the unit square $[0, 1]^2$;

$\langle x, 0 \rangle$ $\langle y, 1 \rangle$		$\langle x, 2 \rangle$ $\langle y, 2 \rangle$
$\langle x, 0 \rangle$ $\langle y, 0 \rangle$	$\langle x, 1 \rangle$ $\langle y, 0 \rangle$	

Figure 4. A strategy guide for **3** that assigns values to x and y

- $p_0 = \langle \phi, H(a, b), \exists \rangle$;
- if $p_i = \langle \sim\psi, s', \alpha \rangle$ then $p_{i+1} = \langle \psi, s', \bar{\alpha} \rangle$, where $\bar{\exists} = \forall$ and $\bar{\forall} = \exists$;
- if $p_i = \langle \psi \vee \chi, s', \alpha \rangle$ then $p_{i+1} = \langle \psi, s', \alpha \rangle$ or $p_{i+1} = \langle \chi, s', \alpha \rangle$;
- if $p_i = \langle \exists x^V \psi, s', \alpha \rangle$ then $p_{i+1} = \langle \psi, s'[x/m], \alpha \rangle$ for some $m \in M$.

We will refer to the real number a as *Eloise's lottery number*, and to b as *Abelard's lottery number*.

- A *terminal history* or *complete play* is a history such that $H(a, b)$ is undefined, in which case neither player wins, or $p_n = \langle \chi, s', \alpha \rangle$, where χ is an atomic formula, in which case
 - player α wins if $\mathbf{M}, s' \models \chi$,
 - player $\bar{\alpha}$ wins if $\mathbf{M}, s' \not\models \chi$.

We simply write $\underline{G}(\mathbf{M}, \phi)$ when H is the strategy guide that sends every point in the unit square to the empty assignment.

A strategy for Eloise in the lottery game $\underline{G}(\mathbf{M}, H, \phi)$ is a function $\underline{\sigma}$ that assigns to every $a \in [0, 1]$ a pure strategy for the semantic game

$$G(\mathbf{M}, H(a, *), \phi).$$

Eloise follows $\underline{\sigma}$ in a history $\langle a, b \rangle, p_0, \dots, p_n$ if $\underline{\sigma}(a)(p_i) = p_{i+1}$ for all $0 \leq i < n$ such that $p_i \in P_{\exists}$. Similarly, a strategy for Abelard in $\underline{G}(\mathbf{M}, H, \phi)$ is

a function $\underline{\tau}$ that assigns to every $b \in [0, 1]$ a pure strategy for

$$G(\mathbf{M}, H(*, b), \phi).$$

Abelard follows $\underline{\tau}$ in $\langle a, b \rangle, p_0, \dots, p_n$ if $\underline{\tau}(b)(p_i) = p_{i+1}$ for all $0 \leq i < n$ such that $p_i \in P_{\forall}$. To avoid confusion, we refer to strategies for lottery games as *lottery-augmented strategies*, in contradistinction to pure strategies and mixed strategies, which will always refer to strategies for semantic games $G(\mathbf{M}, X, \phi)$ as in Definition 2.5.

The lottery numbers a and b allow the players to keep track of which pure strategies they decided to follow at the beginning of the game. It may help the reader's intuition to think of Eloise and Abelard not as single players, but as coalitions of players whose interests are aligned. Before play begins, each coalition decides which pure strategy they are going to follow (possibly with the aid of a coin-flip or the roll of a die). No further coordination between the members of each coalition is allowed during play. However, members of each coalition may be able to observe some—but not necessarily all—of the actions taken by their opponents and their allies before it is their own turn. If so, they are allowed to use their observations when deciding how to act.

For example, imagine that Eloise is a group of CIA agents, and that Abelard is an opposing group of KGB agents. Before leaving headquarters, each group distributes a code book to its members. Once in the field, the agents await a signal telling them which strategy they should execute. For instance, if at a certain time BBC Radio 4 broadcasts an advertisement for a nonexistent brand of laundry detergent, the CIA agents will execute the strategy on page 19 of their code book. At approximately the same time, Soviet state television broadcasts an homage to Vladimir Lenin containing a certain agreed-upon word, telling the KGB agents to execute the strategy on page 42 of their code book. During the course of the operation, a given CIA agent may be able to observe the actions taken by some (but not necessarily all) of her fellow agents. She may also be able to observe certain actions taken by the opposing agents before executing her assignment. Her actions might be observed in turn and affect subsequent choices made by friend and foe alike.

EXAMPLE 5.4. Suppose the CIA is hiding a defector from the Soviet Union in one of three safe houses. The KGB is desperately trying to find the defector, but they only have time to search one of the safe houses before the CIA smuggles her out of the country. The CIA decides which safe house to use according to the flip of a fair three-sided coin, so they tell the defector

$\langle x, 0 \rangle$ $\langle z, 0 \rangle$	$\langle x, 1 \rangle$ $\langle z, 1 \rangle$	$\langle x, 2 \rangle$ $\langle z, 2 \rangle$
--	--	--

Figure 5. Another strategy guide for **3** that assigns values to x and y .

that she has a $2/3$ chance of making it to freedom. Unfortunately, the KGB has a mole inside the CIA who signals the location of the defector to his comrades. The situation can be modeled by the DF formula,

$$\forall y^z x \neq y,$$

viewed as a subformula of the sentence

$$\exists x^0 \forall z^x \forall y^z x \neq y,$$

where x is the location of the defector, y is the safe house searched by the KGB, and z is the signal sent by the mole. The state of the game after the first two moves is encoded in the strategy guide H depicted in Figure 5. Because the value of z signals the location of the defector, the defector is doomed if the KGB follows the obvious strategy defined by $\tau(b) = (y := z)$ for all $b \in [0, 1]$.

Let H be a strategy guide for \mathbf{M} that assigns values to (at least) the free variables of ϕ . For a fixed pair of lottery numbers $a, b \in [0, 1]$, let σ be a pure strategy for Eloise in the semantic game

$$G_{\langle a, * \rangle} = G(\mathbf{M}, H(a, *), \phi),$$

and let τ be a pure strategy for Abelard in

$$G_{\langle *, b \rangle} = G(\mathbf{M}, H(*, b), \phi).$$

If $\langle a, b \rangle \in \text{dom}(H)$ there is a unique terminal history p_0, \dots, p_n of both games in which $p_0 = \langle \phi, H(a, b), \exists \rangle$, Eloise follows σ , and Abelard follows τ . Notice that σ and τ are both strategies for

$$G_{\langle a, b \rangle} = G(\mathbf{M}, \{H(a, b)\}, \phi),$$

and p_0, \dots, p_n is the unique terminal history of $G_{\langle a, b \rangle}$ in which Eloise follows σ and Abelard follows τ .

Thus, given a pair of strategies $\langle \underline{\sigma}, \underline{\tau} \rangle$ for the lottery game $\underline{G}(\mathbf{M}, H, \phi)$ we can define a function

$$\underline{u}(\underline{\sigma}, \underline{\tau}): \text{dom}(H) \rightarrow \{0, 1\}$$

such that $\underline{u}(\underline{\sigma}, \underline{\tau})(a, b) = 1$ if Eloise wins the only complete play of $G_{\langle a, b \rangle}$ in which she follows $\underline{\sigma}(a)$ and Abelard follows $\underline{\tau}(b)$; otherwise $\underline{u}(\underline{\sigma}, \underline{\tau})(a, b) = 0$. We will use the function $\underline{u}(\underline{\sigma}, \underline{\tau})$ to calculate the probability that Eloise wins the lottery game $\underline{G}(\mathbf{M}, H, \phi)$ when she follows $\underline{\sigma}$ and Abelard follows $\underline{\tau}$. In order for the probability that Eloise wins to be well defined, the lottery-augmented strategies $\underline{\sigma}$ and $\underline{\tau}$ must satisfy certain measurability conditions.

Two lottery numbers $a, a' \in [0, 1]$ are said to be *equivalent* with respect to a lottery-augmented strategy $\underline{\sigma}$ if $\underline{\sigma}(a) = \underline{\sigma}(a')$. The equivalence class of a lottery number a with respect to $\underline{\sigma}$ is

$$\llbracket a \rrbracket_{\underline{\sigma}} = \{ a' \in [0, 1] : \underline{\sigma}(a) = \underline{\sigma}(a') \}.$$

A lottery-augmented strategy $\underline{\sigma}$ is *measurable* if $\llbracket a \rrbracket_{\underline{\sigma}}$ is measurable for all $a \in [0, 1]$.

DEFINITION 5.5. Suppose $\langle \underline{\sigma}, \underline{\tau} \rangle$ is a pair of measurable strategies for the lottery game $\underline{G}(\mathbf{M}, H, \phi)$. Then the *expected value* of $\langle \underline{\sigma}, \underline{\tau} \rangle$ is

$$\underline{U}(\underline{\sigma}, \underline{\tau}) = \int_{\text{dom}(H)} \underline{u}(\underline{\sigma}, \underline{\tau})(a, b) da db.$$

Equilibria for lottery games are defined in the obvious way. A pair $\langle \underline{\sigma}^*, \underline{\tau}^* \rangle$ of measurable lottery-augmented strategies is a *lottery equilibrium* if for every pair $\langle \underline{\sigma}, \underline{\tau} \rangle$ of measurable lottery-augmented strategies,

$$\underline{U}(\underline{\sigma}, \underline{\tau}^*) \leq \underline{U}(\underline{\sigma}^*, \underline{\tau}^*) \leq \underline{U}(\underline{\sigma}^*, \underline{\tau}).$$

If $\langle \underline{\sigma}, \underline{\tau} \rangle$ and $\langle \underline{\sigma}', \underline{\tau}' \rangle$ are both lottery equilibria for the same game, then

$$\underline{U}(\underline{\sigma}, \underline{\tau}) = \underline{U}(\underline{\sigma}', \underline{\tau}')$$

for reasons analogous to Proposition 4.4. If the lottery game $\underline{G}(\mathbf{M}, H, \phi)$ has an equilibrium $\langle \underline{\sigma}^*, \underline{\tau}^* \rangle$ we define

$$\underline{\text{Val}}(\mathbf{M}, H, \phi) = \underline{U}(\underline{\sigma}^*, \underline{\tau}^*).$$

We write $\mathbf{M} \models_H^\varepsilon \phi$ if and only if $\underline{\text{Val}}(\mathbf{M}, H, \phi) = \varepsilon$. When H is the strategy guide that sends every point in the unit square to the empty assignment, we simply write $\underline{\text{Val}}(\mathbf{M}, \phi) = \varepsilon$ and $\mathbf{M} \models^\varepsilon \phi$.

The purpose of lottery games is to allow us to simulate the strategic version of a semantic game using an extensive game, which can then be decomposed into its subgames. In order to verify that the simulation is accurate, we must formalize the connection between mixed strategies and lottery-augmented strategies.

DEFINITION 5.6. Let ϕ be a DF sentence, and let \mathbf{M} be a suitable finite structure. Given a mixed strategy μ for Eloise in the semantic game $G(\mathbf{M}, \phi)$ and an enumeration $\sigma_1, \dots, \sigma_m$ of the pure strategies in the support of μ , let $\text{Lot}(\mu)$ denote the measurable strategy for $\underline{G}(\mathbf{M}, \phi)$ defined by $\text{Lot}(\mu)(a) = \sigma_i$ if $a \in A_i$, where

$$\begin{aligned} A_1 &= [0, \mu(\sigma_1)), \\ A_2 &= [\mu(\sigma_1), \mu(\sigma_1) + \mu(\sigma_2)), \\ &\vdots \\ A_m &= \left[\sum_{i=1}^{m-1} \mu(\sigma_i), 1 \right]. \end{aligned}$$

Thus every mixed strategy can be represented by a lottery-augmented strategy. Moreover, every pair of mixed strategies can be represented by a pair of lottery-augmented strategies that yield the same expected utility.

PROPOSITION 5.7. *Let ϕ be a DF sentence, and let \mathbf{M} be a suitable finite structure. If $\langle \mu, \nu \rangle$ is a pair of mixed strategies for $G(\mathbf{M}, \phi)$, then*

$$U(\mu, \nu) = \underline{U}(\text{Lot}(\mu), \text{Lot}(\nu)).$$

PROOF. Let $\sigma_1, \dots, \sigma_m$ enumerate the pure strategies in the support of μ , and for each $1 \leq i \leq m$, let A_i be the corresponding interval of length $\mu(\sigma_i)$ as in Definition 5.6. Similarly, let τ_1, \dots, τ_n enumerate the pure strategies in the support of ν , and for each $1 \leq j \leq n$, let B_j be the corresponding

interval of length $\nu(\tau_j)$. Then

$$\begin{aligned} \underline{U}(\text{Lot}(\mu), \text{Lot}(\nu)) &= \int_0^1 \int_0^1 \underline{u}(\text{Lot}(\mu), \text{Lot}(\nu))(a, b) da db \\ &= \sum_{i=1}^m \sum_{j=1}^n \int_{A_i \times B_j} u(\sigma_i, \tau_j) da db \\ &= \sum_{i=1}^m \sum_{j=1}^n \mu(\sigma_i) \nu(\tau_j) u(\sigma_i, \tau_j). \quad \square \end{aligned}$$

Conversely, every measurable lottery-augmented strategy can be thought of as a mixed strategy, and every pair of measurable lottery-augmented strategies yields the same expected utility as a pair of mixed strategies.

DEFINITION 5.8. Given a measurable strategy $\underline{\sigma}$ for $\underline{G}(\mathbf{M}, \phi)$, let $\text{Mix}(\underline{\sigma})$ denote the mixed strategy for $G(\mathbf{M}, \phi)$ defined by

$$\text{Mix}(\underline{\sigma})(\sigma) = \int_{\underline{\sigma}^{-1}(\sigma)} da$$

for every pure strategy σ for $G(\mathbf{M}, \phi)$. That is, $\text{Mix}(\underline{\sigma})$ tells the appropriate player to follow each pure strategy σ with probability equal to the measure of the subset of $[0, 1]$ that $\underline{\sigma}$ maps to σ .

PROPOSITION 5.9. If $\langle \underline{\sigma}, \underline{\tau} \rangle$ is a pair of measurable strategies for $\underline{G}(\mathbf{M}, \phi)$,

$$\underline{U}(\underline{\sigma}, \underline{\tau}) = U(\text{Mix}(\underline{\sigma}), \text{Mix}(\underline{\tau})).$$

PROOF. Let $\sigma_1, \dots, \sigma_m$ enumerate the pure strategies in the range of $\underline{\sigma}$, and let $A_i = \underline{\sigma}^{-1}(\sigma_i)$. Similarly, let τ_1, \dots, τ_n enumerate the pure strategies in the range of $\underline{\tau}$, and let $B_j = \underline{\tau}^{-1}(\tau_j)$. Then

$$\begin{aligned} U(\text{Mix}(\underline{\sigma}), \text{Mix}(\underline{\tau})) &= \sum_{i=1}^m \sum_{j=1}^n \text{Mix}(\underline{\sigma})(\sigma_i) \text{Mix}(\underline{\tau})(\tau_j) u(\sigma_i, \tau_j) \\ &= \sum_{i=1}^m \sum_{j=1}^n \int_{A_i \times B_j} u(\sigma_i, \tau_j) da db \\ &= \int_0^1 \int_0^1 \underline{u}(\underline{\sigma}, \underline{\tau})(a, b) da db. \quad \square \end{aligned}$$

Strictly speaking, not every measurable lottery-augmented strategy $\underline{\sigma}$ is the representation of a mixed strategy because the preimage of a pure

strategy under $\underline{\sigma}$ may not be an interval. Also, a given mixed strategy μ may have several representations depending on the enumeration of the pure strategies in its support. However, we can easily recover the original mixed strategy from any of its representations.

LEMMA 5.10. *If μ is a mixed strategy for $G(\mathbf{M}, \phi)$, then $\mu = \text{Mix}(\text{Lot}(\mu))$.*

PROOF. For every pure strategy σ ,

$$\text{Mix}(\text{Lot}(\mu))(\sigma) = \int_{\text{Lot}(\mu)^{-1}(\sigma)} da = \mu(\sigma). \quad \square$$

THEOREM 5.11. *Let ϕ be a DF sentence, and let \mathbf{M} be a suitable finite structure. If $\langle \mu^*, \nu^* \rangle$ is a Nash equilibrium for the semantic game $G(\mathbf{M}, \phi)$, then $\langle \text{Lot}(\mu^*), \text{Lot}(\nu^*) \rangle$ is a lottery equilibrium for $\underline{G}(\mathbf{M}, \phi)$, and*

$$\text{Val}(\mathbf{M}, \phi) = \underline{\text{Val}}(\mathbf{M}, \phi).$$

PROOF. Suppose $\langle \mu^*, \nu^* \rangle$ is a Nash equilibrium for $G(\mathbf{M}, \phi)$. Then by Proposition 5.7,

$$\text{Val}(\mathbf{M}, \phi) = U(\mu^*, \nu^*) = \underline{U}(\text{Lot}(\mu^*), \text{Lot}(\nu^*)).$$

To show that $\langle \text{Lot}(\mu^*), \text{Lot}(\nu^*) \rangle$ is a lottery equilibrium, let $\langle \underline{\sigma}, \underline{\tau} \rangle$ be a pair of lottery-augmented strategies for $\underline{G}(\mathbf{M}, \phi)$. Then by Proposition 5.7, Proposition 5.9, and Lemma 5.10,

$$\underline{U}(\underline{\sigma}, \text{Lot}(\nu^*)) = U(\text{Mix}(\underline{\sigma}), \nu^*) \leq U(\mu^*, \nu^*) = \underline{U}(\text{Lot}(\mu^*), \text{Lot}(\nu^*)).$$

Similarly,

$$\underline{U}(\text{Lot}(\mu^*), \underline{\tau}) = U(\mu^*, \text{Mix}(\underline{\tau})) \geq U(\mu^*, \nu^*) = \underline{U}(\text{Lot}(\mu^*), \text{Lot}(\nu^*)). \quad \square$$

6. Lottery semantics

Using lottery games we can formulate a compositional semantics for DF logic that is equivalent to equilibrium semantics on DF sentences. To do so, we need a way to update our strategy guides to reflect the new state of the game after each move.

There are three kinds of moves: negation, disjunction, and existential quantification. Negation tells the players to switch roles, and we can encode such role-reversals by transposing strategy guides.

DEFINITION 6.1. Let H be a strategy guide for \mathbf{M} that assigns values to the variables in V . The *transposition* of H is the function

$$H^T: \text{dom}(H)^{-1} \rightarrow M^V$$

defined by $H^T(a, b) = H(b, a)$.

PROPOSITION 6.2. H^T is a strategy guide.

PROOF. If H respects the grid $\langle \mathcal{A}, \mathcal{B} \rangle$, then H^T respects the grid $\langle \mathcal{B}, \mathcal{A} \rangle$. \square

PROPOSITION 6.3. If $\langle \underline{\sigma}, \underline{\tau} \rangle$ is a pair of strategies for $\underline{G}(\mathbf{M}, H, \psi)$, then we can define strategies $\underline{\sigma}^T$ and $\underline{\tau}^T$ for $\underline{G}(\mathbf{M}, H^T, \sim\psi)$ by

$$\begin{aligned} \underline{\tau}^T(a)(\chi, s, \exists) &= \underline{\tau}(a)(\chi, s, \forall), \\ \underline{\sigma}^T(b)(\chi, s, \forall) &= \underline{\sigma}(b)(\chi, s, \exists), \end{aligned}$$

respectively, and $\underline{U}(\underline{\tau}^T, \underline{\sigma}^T) = \int_{\text{dom}(H)} da db - \underline{U}(\underline{\sigma}, \underline{\tau})$.

PROOF. Every play $\langle a, b \rangle, p_0, \dots, p_n$ of $\underline{G}(\mathbf{M}, H, \psi)$ in which Eloise follows $\underline{\sigma}$ and Abelard follows $\underline{\tau}$ corresponds to a play

$$\langle b, a \rangle, \langle \sim\psi, H(a, b), \exists \rangle, \bar{p}_0, \dots, \bar{p}_n$$

of $\underline{G}(\mathbf{M}, H^T, \sim\psi)$, where

$$p_i = \langle \chi, s, \alpha \rangle \quad \text{implies} \quad \bar{p}_i = \langle \chi, s, \bar{\alpha} \rangle.$$

In particular, every terminal history $\langle a, b \rangle, p_0, \dots, p_n$ of $\underline{G}(\mathbf{M}, H, \psi)$ with $p_n = \langle \chi, s, \alpha \rangle$ corresponds to a terminal history

$$\langle b, a \rangle, \langle \sim\psi, H(a, b), \exists \rangle, \bar{p}_0, \dots, \bar{p}_n$$

of $\underline{G}(\mathbf{M}, H^T, \sim\psi)$ with $\bar{p}_n = \langle \chi, s, \bar{\alpha} \rangle$, and vice versa. Thus, for all $\langle a, b \rangle \in \text{dom}(H)$,

$$\underline{u}(\underline{\tau}^T, \underline{\sigma}^T)(b, a) = 1 - \underline{u}(\underline{\sigma}, \underline{\tau})(a, b).$$

Therefore,

$$\begin{aligned} \underline{U}(\underline{\tau}^T, \underline{\sigma}^T) &= \int_{\text{dom}(H^T)} \underline{u}(\underline{\tau}^T, \underline{\sigma}^T)(b, a) db da \\ &= \int_{\text{dom}(H)} 1 - \underline{u}(\underline{\sigma}, \underline{\tau})(a, b) da db \\ &= \int_{\text{dom}(H)} da db - \underline{U}(\underline{\sigma}, \underline{\tau}). \end{aligned} \quad \square$$

A disjunction prompts the verifier to choose the left or right disjunct. If he or she chooses the left disjunct, what might have happened had he or she chosen the right disjunct does not affect the outcome of the current play. Hence we can split the strategy guide H into two subguides, and consider each case separately.

DEFINITION 6.4. Let ϕ be a DF formula, \mathbf{M} a finite structure, and V be a finite set of variables that contains $\text{Free}(\phi)$. A *splitting function for \mathbf{M} over V* assigns to each $a \in [0, 1]$ a function

$$\text{Sp}_a: M^V \rightarrow \{\text{left}, \text{right}\}.$$

A splitting function Sp is *measurable* if for all $a \in [0, 1]$, the equivalence class

$$\llbracket a \rrbracket_{\text{Sp}} = \{ a' \in [0, 1] : \text{Sp}_a = \text{Sp}_{a'} \}$$

is measurable.

If H is a strategy guide for \mathbf{M} that assigns values to the variables in V , we say that a splitting function Sp for \mathbf{M} over V *splits H* into two functions H_1 and H_2 such that

$$\begin{aligned} \text{dom}(H_1) &= \text{dom}(H) \cap \left\{ \langle a, b \rangle \in [0, 1]^2 : \text{Sp}_a(H(a, b)) = \text{left} \right\}, \\ \text{dom}(H_2) &= \text{dom}(H) \cap \left\{ \langle a, b \rangle \in [0, 1]^2 : \text{Sp}_a(H(a, b)) = \text{right} \right\}, \end{aligned}$$

and H_1, H_2 both agree with H on their respective domains.

PROPOSITION 6.5. *If Sp is measurable, then H_1 and H_2 are strategy guides.*

PROOF. Suppose H respects the grid $\langle \mathcal{A}, \mathcal{B} \rangle$. If Sp is measurable, then every cell in the partition

$$\mathcal{A}' = \{ A \cap \llbracket a \rrbracket_{\text{Sp}} : A \in \mathcal{A}, a \in [0, 1] \}$$

is measurable. Hence $\langle \mathcal{A}', \mathcal{B} \rangle$ is a grid.

To show that H_1 and H_2 respect $\langle \mathcal{A}', \mathcal{B} \rangle$, suppose $a, a' \in A \cap \llbracket a \rrbracket_{\text{Sp}} \in \mathcal{A}'$, where $A \in \mathcal{A}$, and $b, b' \in B \in \mathcal{B}$. Then $H(a, b) = H(a', b')$ (or both are undefined) because H respects $\langle \mathcal{A}, \mathcal{B} \rangle$. Furthermore, $\text{Sp}_a = \text{Sp}_{a'}$ because $a' \in \llbracket a \rrbracket_{\text{Sp}}$, so if $H(a, b)$ and $H(a', b')$ are defined,

$$\text{Sp}_a(H(a, b)) = \text{Sp}_{a'}(H(a', b')).$$

Hence both $\langle a, b \rangle$ and $\langle a', b' \rangle$ belong to $\text{dom}(H_1)$, or both belong to $\text{dom}(H_2)$. In either case $H_i(a, b) = H_i(a', b')$. \square

An existential quantifier prompts the verifier to choose the value of the quantified variable. Thus we need a way to extend (“supplement”) the current assignment.

DEFINITION 6.6. Let ϕ be a DF formula, \mathbf{M} a finite structure, and V a finite set of variables that contains $\text{Free}(\phi)$. A *supplementation function* for \mathbf{M} over $W \subseteq V$ assigns to every $a \in [0, 1]$ a function

$$F_a: M^W \rightarrow M.$$

A supplementation function is *measurable* if, for every $a \in [0, 1]$, the equivalence class

$$\llbracket a \rrbracket_F = \{ a' \in [0, 1] : F_a = F_{a'} \}$$

is measurable.

If H is a strategy guide for \mathbf{M} that assigns values to the variables in V , we say that a supplementation function F for \mathbf{M} over $W \subseteq V$ *supplements* H by assigning a value to x , yielding a function $H[x/F]$ defined by

$$H[x/F](a, b) = H(a, b) \left[x / F_a(H(a, b)) \right].$$

PROPOSITION 6.7. *If F is measurable, then $H[x/F]$ is a strategy guide.*

PROOF. Suppose H respects the grid $\langle \mathcal{A}, \mathcal{B} \rangle$. If F is measurable, then every cell in the partition

$$\mathcal{A}' = \{ A \cap \llbracket a \rrbracket_F : a \in [0, 1] \}$$

is measurable. Hence $\langle \mathcal{A}', \mathcal{B} \rangle$ is a grid.

To show that $H[x/F]$ respects $\langle \mathcal{A}', \mathcal{B} \rangle$, suppose $a, a' \in A \cap \llbracket a \rrbracket_F \in \mathcal{A}'$, where $A \in \mathcal{A}$, and $b, b' \in B \in \mathcal{B}$. Then

$$\begin{aligned} H[x/F](a, b) &= H(a, b) \left[x / F_a(H(a, b)) \right] \\ &= H(a', b') \left[x / F_{a'}(H(a', b')) \right] \\ &= H[x/F](a', b'). \end{aligned} \quad \square$$

Now that we know how to update strategy guides to account for the effects of the players’ moves, we are ready to prove our main theorem.

THEOREM 6.8 (Lottery Semantics). *Let ϕ be a DF formula, \mathbf{M} a suitable finite structure, and H a strategy guide for \mathbf{M} that assigns values to the variables in a finite set V that contains $\text{Free}(\phi)$. Suppose $\langle \underline{\sigma}^*, \underline{\tau}^* \rangle$ is an equilibrium for the lottery game $\underline{G}(\mathbf{M}, H, \phi)$.*

- If ϕ is atomic,

$$\underline{\text{Val}}(\mathbf{M}, H, \phi) = \int_{\{\langle a,b \rangle : \mathbf{M} \models_{H(a,b)} \phi\}} da db.$$

- If ϕ is $\sim\psi$,

$$\underline{\text{Val}}(\mathbf{M}, H, \phi) = \int_{\text{dom}(H)} da db - \underline{\text{Val}}(\mathbf{M}, H^T, \psi).$$

- If ϕ is $\psi_1 \vee \psi_2$,

$$\underline{\text{Val}}(\mathbf{M}, H, \phi) = \max_{\text{Sp}} \left(\underline{\text{Val}}(\mathbf{M}, H_1, \psi_1) + \underline{\text{Val}}(\mathbf{M}, H_2, \psi_2) \right),$$

where the maximum is taken over all measurable splitting functions Sp for \mathbf{M} over V , and H_1, H_2 are the strategy guides into which Sp splits H .

- If ϕ is $\exists x^W \psi$,

$$\underline{\text{Val}}(\mathbf{M}, H, \phi) = \max_F \underline{\text{Val}}(\mathbf{M}, H[x/F], \psi),$$

where the maximum is taken over all measurable supplementation functions F for \mathbf{M} over W .

PROOF. If ϕ is atomic, every play of the lottery game $\underline{G}(\mathbf{M}, H, \phi)$ is complete and has the form $\langle a, b \rangle, \langle \phi, H(a, b), \exists \rangle$, so $\underline{\sigma}^*$ must assign the empty pure strategy to every $a \in [0, 1]$. Similarly, $\underline{\tau}^*$ must assign the empty pure strategy to every $b \in [0, 1]$. Then for all $\langle a, b \rangle \in \text{dom}(H)$,

$$\underline{u}(\underline{\sigma}^*, \underline{\tau}^*)(a, b) = \begin{cases} 1 & \text{if } \mathbf{M} \models_{H(a,b)} \phi, \\ 0 & \text{if } \mathbf{M} \not\models_{H(a,b)} \phi. \end{cases}$$

Hence

$$\begin{aligned} \underline{\text{Val}}(\mathbf{M}, H, \phi) &= \int_{\text{dom}(H)} \underline{u}(\underline{\sigma}^*, \underline{\tau}^*)(a, b) da db \\ &= \int_{\{\langle a,b \rangle : \mathbf{M} \models_{H(a,b)} \phi\}} da db. \end{aligned}$$

Suppose ϕ is $\sim\psi$. Then $\langle \underline{\tau}', \underline{\sigma}' \rangle$ is a lottery equilibrium for $\underline{G}(\mathbf{M}, H^T, \psi)$, where $\underline{\sigma}'$ and $\underline{\tau}'$ are the lottery-augmented strategies induced by $(\underline{\sigma}^*)^T$ and

$(\tau^*)^T$, respectively. Furthermore, by Proposition 6.3,

$$\begin{aligned} \underline{\text{Val}}(\mathbf{M}, H, \sim\psi) &= \int_{\text{dom}(H^T)} db da - \underline{U}(\tau', \sigma') \\ &= \int_{\text{dom}(H)} da db - \underline{\text{Val}}(\mathbf{M}, H^T, \psi). \end{aligned}$$

Suppose ϕ is $\psi_1 \vee \psi_2$. Then the splitting function defined by

$$\text{Sp}_a^*(s) = \begin{cases} \text{left} & \text{if } \underline{\sigma}^*(a)(\phi, s, \exists) = \langle \psi_1, s, \exists \rangle, \\ \text{right} & \text{if } \underline{\sigma}^*(a)(\phi, s, \exists) = \langle \psi_2, s, \exists \rangle. \end{cases}$$

is measurable because for all $a \in [0, 1]$,

$$\llbracket a \rrbracket_{\text{Sp}^*} = \bigcup \left\{ \llbracket a' \rrbracket_{\underline{\sigma}^*} : (\forall s \in M^V) [\underline{\sigma}^*(a)(\phi, s, \exists) = \underline{\sigma}^*(a')(\phi, s, \exists)] \right\}.$$

There are only finitely many $\llbracket a' \rrbracket_{\underline{\sigma}^*}$, each of which is measurable by hypothesis. Thus Sp^* splits H into two strategy guides H_1^* and H_2^* such that

$$\text{dom}(H_i^*) = \left\{ \langle a, b \rangle \in \text{dom}(H) : \underline{\sigma}^*(a)(\phi, H(a, b), \exists) = \langle \psi_i, H(a, b), \exists \rangle \right\}.$$

Let $\underline{\sigma}'$ and $\underline{\tau}'$ be the strategies for $\underline{G}(\mathbf{M}, H_1^*, \psi_1)$ —and let $\underline{\sigma}''$ and $\underline{\tau}''$ be the strategies for $\underline{G}(\mathbf{M}, H_2^*, \psi_2)$ —induced by $\underline{\sigma}^*$ and $\underline{\tau}^*$, respectively. Then $\langle \underline{\sigma}', \underline{\tau}' \rangle$ and $\langle \underline{\sigma}'', \underline{\tau}'' \rangle$ are lottery equilibria for their respective games, else $\langle \underline{\sigma}^*, \underline{\tau}^* \rangle$ would not be an equilibrium for $\underline{G}(\mathbf{M}, H, \phi)$. Moreover,

$$\begin{aligned} \underline{\text{Val}}(\mathbf{M}, H, \psi_1 \vee \psi_2) &= \int_{\text{dom}(H)} \underline{u}(\underline{\sigma}^*, \underline{\tau}^*)(a, b) da db \\ &= \int_{\text{dom}(H_1^*)} \underline{u}(\underline{\sigma}', \underline{\tau}')(a, b) da db \\ &\quad + \int_{\text{dom}(H_2^*)} \underline{u}(\underline{\sigma}'', \underline{\tau}'')(a, b) da db \\ &= \underline{\text{Val}}(\mathbf{M}, H_1^*, \psi_1) + \underline{\text{Val}}(\mathbf{M}, H_2^*, \psi_2). \end{aligned}$$

Now let Sp be a measurable splitting function that splits H into H_1 and H_2 , and suppose $\langle \underline{\sigma}^{(1)}, \underline{\tau}^{(1)} \rangle$ and $\langle \underline{\sigma}^{(2)}, \underline{\tau}^{(2)} \rangle$ are equilibria for the games $\underline{G}(\mathbf{M}, H_1, \psi_1)$ and $\underline{G}(\mathbf{M}, H_2, \psi_2)$, respectively. Define strategies $\underline{\sigma}$ and $\underline{\tau}$ for

$\underline{G}(\mathbf{M}, H, \psi_1 \vee \psi_2)$ by

$$\underline{\sigma}(a)(\psi_1 \vee \psi_2, s, \exists) = \begin{cases} \langle \psi_1, s, \exists \rangle & \text{if } \text{Sp}_a(s) = \text{left}, \\ \langle \psi_2, s, \exists \rangle & \text{if } \text{Sp}_a(s) = \text{right}, \end{cases}$$

$$\underline{\sigma}(a)(\chi, s, \exists) = \begin{cases} \underline{\sigma}^{(1)}(a)(\chi, s, \exists) & \text{if } \chi \text{ is a subformula of } \psi_1, \\ \underline{\sigma}^{(2)}(a)(\chi, s, \exists) & \text{if } \chi \text{ is a subformula of } \psi_2, \end{cases}$$

and

$$\underline{\tau}(b)(\chi, s, \forall) = \begin{cases} \underline{\tau}^{(1)}(b)(\chi, s, \forall) & \text{if } \chi \text{ is a subformula of } \psi_1, \\ \underline{\tau}^{(2)}(b)(\chi, s, \forall) & \text{if } \chi \text{ is a subformula of } \psi_2. \end{cases}$$

Then

$$\begin{aligned} \underline{U}(\underline{\sigma}, \underline{\tau}) &= \int_{\text{dom}(H)} \underline{u}(\underline{\sigma}, \underline{\tau})(a, b) da db \\ &= \int_{\text{dom}(H_1)} \underline{u}(\underline{\sigma}^{(1)}, \underline{\tau}^{(1)})(a, b) da db \\ &\quad + \int_{\text{dom}(H_2)} \underline{u}(\underline{\sigma}^{(2)}, \underline{\tau}^{(2)})(a, b) da db \\ &= \underline{U}(\underline{\sigma}^{(1)}, \underline{\tau}^{(1)}) + \underline{U}(\underline{\sigma}^{(2)}, \underline{\tau}^{(2)}). \end{aligned}$$

The pair $\langle \underline{\sigma}, \underline{\tau} \rangle$ may not be a lottery equilibrium for $\underline{G}(\mathbf{M}, H, \psi_1 \vee \psi_2)$, but we know that Abelard cannot decrease $\underline{U}(\underline{\sigma}, \underline{\tau})$ by modifying $\underline{\tau}$ because, if he could, then either $\langle \underline{\sigma}^{(1)}, \underline{\tau}^{(1)} \rangle$ or $\langle \underline{\sigma}^{(2)}, \underline{\tau}^{(2)} \rangle$ would not be a lottery equilibrium. Therefore,

$$\begin{aligned} \underline{\text{Val}}(\mathbf{M}, H_1, \psi_1) + \underline{\text{Val}}(\mathbf{M}, H_2, \psi_2) &= \underline{U}(\underline{\sigma}, \underline{\tau}) \\ &\leq \underline{U}(\underline{\sigma}, \underline{\tau}^*) \\ &\leq \underline{U}(\underline{\sigma}^*, \underline{\tau}^*) \\ &= \underline{\text{Val}}(\mathbf{M}, H, \psi_1 \vee \psi_2). \end{aligned}$$

Suppose ϕ is $\exists x^W \psi$. Let F^* be the measurable supplementation function for \mathbf{M} over W such that

$$\underline{\sigma}^*(a)(\exists x^W \psi, s, \exists) = \langle \psi, s[x/F_a^*(s)], \exists \rangle,$$

and consider the strategies $\underline{\sigma}'$ and $\underline{\tau}'$ for $\underline{G}(\mathbf{M}, H[x/F^*], \psi)$ induced by $\underline{\sigma}^*$ and $\underline{\tau}^*$, respectively. Then $\underline{u}(\underline{\sigma}^*, \underline{\tau}^*) = \underline{u}(\underline{\sigma}', \underline{\tau}')$; hence

$$\underline{\text{Val}}(\mathbf{M}, H, \exists x^V \psi) = \underline{U}(\underline{\sigma}', \underline{\tau}').$$

Indeed $\langle \underline{\sigma}', \underline{\tau}' \rangle$ is an equilibrium for $\underline{G}(\mathbf{M}, H[x/F^*], \psi)$ because if either player could improve his or her expected utility in this game, they could also improve their expected utility in $\underline{G}(\mathbf{M}, H, \exists x^W \psi)$. Thus

$$\underline{\text{Val}}(\mathbf{M}, H, \exists x^V \psi) = \underline{\text{Val}}(\mathbf{M}, H[x/F^*], \psi).$$

Now let F be a measurable supplementation function for \mathbf{M} over W , and suppose $\langle \underline{\sigma}'', \underline{\tau}'' \rangle$ is a lottery equilibrium for $\underline{G}(\mathbf{M}, H[x/F], \psi)$. Extend $\underline{\sigma}''$ and $\underline{\tau}''$ to strategies $\underline{\sigma}$ and $\underline{\tau}$ for $\underline{G}(\mathbf{M}, H, \exists x^W \psi)$ by setting

$$\begin{aligned} \underline{\sigma}(a)(\exists x^W \psi, s, \exists) &= \langle \psi, s[x/F_a(s)], \exists \rangle, \\ \underline{\sigma}(a)(\chi, s, \exists) &= \underline{\sigma}''(\chi, s, \exists), \end{aligned}$$

and

$$\underline{\tau}(\chi, s, \forall) = \underline{\tau}''(\chi, s, \forall).$$

Then $\underline{u}(\underline{\sigma}, \underline{\tau}) = \underline{u}(\underline{\sigma}'', \underline{\tau}'')$, and hence $\underline{U}(\underline{\sigma}, \underline{\tau}) = \underline{U}(\underline{\sigma}'', \underline{\tau}'')$. The pair $\langle \underline{\sigma}, \underline{\tau} \rangle$ may not be a lottery equilibrium for $\underline{G}(\mathbf{M}, H, \exists x^W \psi)$, but we know that Abelard cannot decrease $\underline{U}(\underline{\sigma}, \underline{\tau})$ by modifying $\underline{\tau}$ because, if he could, then $\langle \underline{\sigma}'', \underline{\tau}'' \rangle$ would not be a lottery equilibrium for $\underline{G}(\mathbf{M}, H[x/F], \psi)$. Therefore,

$$\begin{aligned} \underline{\text{Val}}(\mathbf{M}, H[x/F], \psi) &= \underline{U}(\underline{\sigma}'', \underline{\tau}'') \\ &= \underline{U}(\underline{\sigma}, \underline{\tau}) \\ &\leq \underline{U}(\underline{\sigma}, \underline{\tau}^*) \\ &\leq \underline{U}(\underline{\sigma}^*, \underline{\tau}^*) \\ &= \underline{\text{Val}}(\mathbf{M}, H, \exists x^W \psi). \quad \square \end{aligned}$$

Conclusion

The connection between game-theoretic semantics and trump semantics is central to the study of first-order logic with imperfect information. Equilibrium semantics is a more refined version of game-theoretic semantics that allows us to assign intermediate truth values to undetermined sentences in finite models. It is our hope that lottery semantics — which bears the same relation to equilibrium semantics as trump semantics does to the usual game-theoretic semantics — will shed similar light on this intriguing new framework.

Several colleagues have noticed the similarity between lottery semantics and fuzzy logic. We thank an anonymous referee for pointing out recent work by Cintula and Majer [2] on the game-theoretic semantics of fuzzy logic. We are eager to analyze the connection between these two different approaches to logic with imperfect information.

Acknowledgements. Both authors wish to gratefully acknowledge the support of the European Science Foundation EUROCORES program Log-ICCC [FP002 - Logic for Interaction (LINT)]. Funding for the second author was provided by the Academy of Finland (grant 129208).

References

- [1] BLASS, A., and Y. GUREVICH, ‘Henkin quantifiers and complete problems’, *Annals of Pure and Applied Logic*, 32 (1986), 1–16.
- [2] CINTULA, P., and O. MAJER, ‘Towards evaluation games for fuzzy logics’, in O. Majer, A.-V. Pietarinen, and T. Tulenheimo, (eds.), *Games: Unifying Logic, Language, and Philosophy*, vol. 15 of *Logic, Epistemology, and the Unity of Science*, Springer, 2009, pp. 117–138.
- [3] ENDERTON, H.B., ‘Finite partially-ordered quantifiers’, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 16 (1970), 393–397.
- [4] GALE, D., and F.M. STEWART, ‘Infinite games with perfect information’, in H.W. Kuhn, and A.W. Tucker, (eds.), *Contributions to the Theory of Games*, vol. 2, Princeton University Press, Princeton, 1953, pp. 245–266.
- [5] GALLIANI, P., *Game Values and Equilibria for Undetermined Sentences of Dependence Logic*, MSc thesis, ILLC Publications, MoL-2008-08, 2008.
- [6] GALLIANI, P., *Probabilistic Dependence Logic*, Forthcoming 2011.
- [7] HENKIN, L., ‘Some remarks on infinitely long formulas’, in *Infinitistic Methods. Proc. Symposium on Foundations of Mathematics*, Pergamon Press, Oxford, 1961, pp. 167–183.
- [8] HINTIKKA, J., ‘Language-games for quantifiers’, in *Studies in Logical Theory*, vol. 2 of *American Philosophical Quarterly Monograph Series*, Basil Blackwell, Oxford, 1968, pp. 46–72.
- [9] HINTIKKA, J., ‘Quantifiers vs. quantification theory’, *Dialectica*, 27 (1973), 329–358.
- [10] HINTIKKA, J., *The Principles of Mathematics Revisited*, Cambridge University Press, Cambridge, 1996.
- [11] HINTIKKA, J., and J. KULAS, *The Game of Language: Studies in Game-Theoretical Semantics and Its Applications*, D. Reidel Publishing Company, Dordrecht, 1983.
- [12] HINTIKKA, J., and G. SANDU, ‘Informational independence as a semantic phenomenon’, in J. E. Fenstad, I. T. Frolov, and R. Hilpinen, (eds.), *Logic, Methodology and Philosophy of Science VIII*, Elsevier, Amsterdam, 1989, pp. 571–589.
- [13] HINTIKKA, J., and G. SANDU, ‘Game-theoretical semantics’, in J. van Benthem, and A. ter Meulen, (eds.), *Handbook of Logic and Language*, Elsevier, Amsterdam, 1997, pp. 361–410.
- [14] HODGES, W., ‘Compositional semantics for a language of imperfect information’, *Logic Journal of IGPL*, 5 (1997), 539–563.
- [15] HODGES, W., ‘Some strange quantifiers’, in J. Mycielski, G. Rozenberg, and A. Salomaa, (eds.), *Structures in Logic and Computer Science: A Selection of Essays in Honor of A. Ehrenfeucht*, no. 1261 in *Lecture Notes in Computer Science*, Springer, 1997, pp. 51–65.

- [16] KUHN, H.W., ‘Extensive games and the problem of information’, in *Contributions to the Theory of Games III*, Princeton University Press, Princeton, 1953, pp. 193–216.
- [17] MANN, A.L., G. SANDU, and M. SEVENSTER, *Independence-Friendly Logic: A Game-Theoretic Approach*, no. 386 in London Mathematical Society Lecture Note Series, Cambridge University Press, Cambridge, 2011.
- [18] PEIRCE, C.S., *Reasoning and the Logic of Things: The Cambridge Conferences Lectures of 1898*, Harvard Historical Studies, Harvard University Press, Cambridge, Massachusetts, 1992.
- [19] SEVENSTER, M., *Branches of Imperfect Information: Logic, Games, and Computation*, Ph.D. thesis, Institute for Logic, Language and Computation, Universiteit van Amsterdam, 2006.
- [20] SEVENSTER, M., and G. SANDU, ‘Equilibrium semantics of languages of imperfect information’, *Annals of Pure and Applied Logic*, 161 (2010), 618–631.
- [21] VÄÄNÄNEN, J., *Dependence Logic: A New Approach to Independence Friendly Logic*, no. 70 in London Mathematical Society Student Texts, Cambridge University Press, Cambridge, 2007.
- [22] VON NEUMANN, J., ‘Zur Theorie der Gesellschaftsspiele’, *Mathematische Annalen*, 100 (1928), 295–320.
- [23] WALKOE, W.J., JR., ‘Finite partially-ordered quantification’, *Journal of Symbolic Logic*, 35 (1970), 535–555.

PIETRO GALLIANI
FNWI
ILLC
Universiteit van Amsterdam
P.O. Box 94242
1090 GE Amsterdam
The Netherlands
pgallian@gmail.com

ALLEN L. MANN
Department of Mathematics
Colgate University
13 Oak Drive
Hamilton, NY 13346
USA
allen.l.mann@gmail.com