THE JOURNAL OF SYMBOLIC LOGIC
Volume 00, Number 0, XXX 0000

THERE ARE ONLY 2 PATTERNS OF SEMANTIC PARADOXES

MICHAŁ WALICKI

Abstract. Using a graph representation of classical logic, the paper shows that the liar or Yablo pattern occurs in every semantic paradox. The core graph theoretic result generalizes theorem of Richardson, showing solvability of finite graphs without odd cycles, to arbitrary graphs which are proven solvable when no odd cycles nor patterns generalizing Yablo's occur. This follows from an earlier result by a new compactness-like theorem, holding for infinitary logic and utilizing the graph representation.

§1. Introduction. It had long seemed that all paradoxes result from a kind of vicious circle, until Yablo produced a non-circular one. Without arguing about its (non-)circularity, we show that a version of the liar or Yablo occurs in every semantic paradox. This intuitive claim needs more precision, as both vicious circularity and Yablo pattern are rather vague notions. Vicious circle may be an impredicative definition as much as liar's self-negation. Using graphs for representing logical theories [1,3,6,8], and eventually just their syntax [2,4,11], allows to define these concepts in the language of graph theory and to apply its results. Vicious circles, as general liar patterns, are then just odd cycles, while even ones capture innocent forms of self-reference. With this reading, Richardson's theorem from the 1950-ties, stating solvability of finite graphs without odd cycles, implies that each finite paradox involves a vicious circle [9].

Unlike the liar, captured by odd cycles of negations, Yablo pattern is more loose and elusive. The basic Yablo paradox can be a mere unwinding of the liar [4], but it is easy to envison a wide range of forms not arising from any finite counterparts [10,14]. In [1], exclusion of cycles and of a general Yablo pattern, formulated as a graph minor condition, is conjectured to avoid paradoxes. (Interpretation requires there exclusion of all, not only odd, cycles.) A different formulation of the same pattern was given independently in [12], along with a proof that excluding it and odd cycles suffices for solvability of graphs with finitely many ends. (An end of a graph is a way along which its paths converge towards infinity, e.g., Yablo graph has only one.) The paper shows this for all graphs, lifting the result for finitely many ends by means of a compactness-like theorem: a graph is solvable if each of its induced subgraphs with finitely many ends is. Its translation into logical language yields a form of compactness for infinitary logic.

§2. Preliminaries. As in [2,4,11] graphs represent (syntax of) logical theories, with sentences as vertices, edges marking negations of targets, and multiple out-neighbours representing conjunction. Every theory in propositional, also infinitary, classical logic has an equivalent one in graph normal form, GNF, giving rise to such a representation [2]. (Related graph representation captures first or higher-order languages [13].) A propositional formula in GNF has the form

$$x \leftrightarrow \bigwedge_{y \in I_x} \neg y, \tag{2.1}$$

where x, y are atoms (propositional variables); $I_x = \emptyset$ yields x, i.e., $x \leftrightarrow \top$. A theory is in GNF when all its formulas are in GNF and every atom occurs exactly once unnegated (on the left of

0022-4812/00/0000-0000/$00.00

$\leftrightarrow$) in such a formula.[1] A paradox is identified with an inconsitent theory in GNF, even though every theory can be written in this form, so conditions for its (in)consistency are not limited to paradoxes. Yet, GNF provides a natural representation of circularity and paradoxes [4, 11].

EXAMPLE 2.2. *GNFs represent naturally T-biconditionals defining truth of the atoms on left sides. Yablo paradox becomes* $Y = \{x_i \leftrightarrow \bigwedge_{i<j<\omega} \neg x_j \mid i \in \omega\}$*, the liar* $\{x \leftrightarrow \neg x\}$*, and example (a)-(d) below is the theory in the middle:*

(a) *This and the next statement are false.* $a \leftrightarrow \neg b \wedge \neg a$

(b) *The next statement is false.* $b \leftrightarrow \neg c$

(c) *The previous statement is false.* $c \leftrightarrow \neg b$

(d) *This and the previous statement are false.* $d \leftrightarrow \neg d \wedge \neg c$

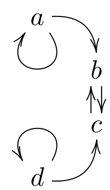

A theory in GNF gives a graph with atoms as vertices and edges from each atom x to all atoms in I_x, the right side of its equivalence (2.1), like the graph to the right above. Graph of Y is $(\omega, <)$, while of the liar $(\{x\}, \{(x,x)\})$, i.e., $x \circlearrowleft$. By "graph" we mean a digraph $G = (\mathbf{V}_G, \mathbf{E}_G)$, $\mathbf{E}_G \subseteq \mathbf{V}_G \times \mathbf{V}_G$, with subscript dropped when an arbitrary or fixed graph is meant. Conversely, each graph gives a theory with axiom $x \leftrightarrow \bigwedge_{y\in\mathbf{E}(x)} \neg y$, for each $x \in \mathbf{V}$.

These syntactic correspondences are reflected by the semantics. Models of GNF theories are in bijection to kernels of the corresponding graphs [2], where *kernel* of a graph is a subset $K \subseteq \mathbf{V}$ such that $K = \mathbf{V} \setminus \mathbf{E}^-(K)$, with $\mathbf{E}^-$ being the converse of $\mathbf{E}$. A graph is *solvable* when it has a kernel. Each theory in Example 2.2 is inconsistent and their graphs have no kernels.

A *sink* (vertex without out-neighbours), belongs to every kernel. Without loss of generality, we limit attention to sinkless graphs, because a graph has a kernel iff its sinkless residuum (obtained by recursively excluding sinks and their predecessors) has one [2]. This amounts simply to propagating the value 'true' of all x with axiom $x \leftrightarrow \top$ through all other axioms.

A *ray* r is a simple outgoing ω path. $\vec{G}$ denotes all rays in G. A *tail* of a ray r is r without some finite prefix. A vertex v *dominates* ray r if it has an infinite set of mutually disjoint (except for the starting v) paths to r. A graph is *safe* if it has no odd cycle and no ray with infinitely many vertices dominating it. The claim from the title, put precisely, says that every safe graph is solvable. Safety is inherited by all subgraphs, so a safe graph is actually *kernel perfect*, KP, that is, every induced subgraph is solvable. (Subgraph of G *induced by* $X \subseteq \mathbf{V}_G$ is $G[X] = (X, (X \times X) \cap \mathbf{E}_G)$.)

The proof uses the fact that a safe graph is solvable when it has finitely many ends [12]. An *end* of a graph is an equivalence class of rays $\{r \in \vec{G} \mid r \simeq r_0\}$ for any $r_0 \in \vec{G}$, where relation $r_1 \simeq r_2$ is the greatest equivalence contained in the preorder $r_1 \preceq r_2$, holding if every tail of r_1 has a path to r_2. Equivalently, $r_1 \preceq r_2$ if $\mathbf{E}^{*}_{-}(r_1) \subseteq \mathbf{E}^{*}_{-}(r_2)$, where $\mathbf{E}^{*}_{-}/\mathbf{E}^{*}$ is reflexive and transitive closure of $\mathbf{E}^-/\mathbf{E}$. An end e determines a subgraph induced by all vertices reaching any of its rays, $\mathbf{E}^{*}_{-}(r_0)$, for any $r_0 \in e$, but such a subgraph may also contain other ends. Preorder $\preceq$ on rays extends to the partial order on ends, $e_1 \preceq e_2$ if $r_1 \preceq r_2$, for any $r_i \in e_i$, or else if $\mathbf{E}^{*}_{-}(r_1) \subseteq \mathbf{E}^{*}_{-}(r_2)$. By $\vec{\vec{G}}$ we denote the set of all ends of graph G, treated often as just described induced subgraphs.

Another central notion to be used generalizes kernels to *semikernels*, namely, subsets $L \subseteq \mathbf{V}$ such that $\mathbf{E}(L) \subseteq \mathbf{E}^-(L) \subseteq \mathbf{V} \setminus L$. Empty semikernel always exists, but we exclude it and saying "semikernel" mean a nonempty one. A graph is KP iff every rem:[nonempty] induced subgraph has a semikernel [7], and our proof establishes this equivalent property.

§3. The proof. The main argument uses a generalization of 'Poison Game' from [5]. For a brief comparison, we quote the original version.

POISON GAME 3.1. *A starts choosing a vertex, and then* B, A *choose alternately vertices from out-neighbours of the last vertex chosen by the opponent. B poisons vertices which he visits, and can re-visit them, but A dies on entering such a vertex, losing the game. A wins if he survives.*

[1] Formula $a \leftrightarrow \neg b$ is in GNF but the theory $\{a \leftrightarrow \neg b\}$ is not, due to the loose b. Such cases can be treated as abbreviations, here, with a fresh atom $\overline{b}$ and two additional formulas $b \leftrightarrow \neg\overline{b}$ and $\overline{b} \leftrightarrow \neg b$.

By Theorem 1 of [5], a finite graph has a semikernel iff A survives a game. On the graph from Example 2.2, A survives starting from any vertex, since reaching c (or b), each player can only keep choosing the same vertex (opposite to its opponent's) *ad infinitum.* Both $\{b\}$ and $\{c\}$ are semikernels. Our version of the game works also for infinite (sinkless) graphs.

Poison Game 3.2. *A starts choosing a vertex, and then B, A choose alternately vertices: A from the out-neighbours of the last vertex chosen by B, while B from the out-neighbours of all vertices chosen so far by A. A poisons all in- and out-neighbours of its choices. A loses the game visiting a poisoned vertex (B can do it unharmed), and wins if he survives.*

While A must choose an out-neighbour of the last vertex chosen by B, B can now jump to out-neighbours of vertices chosen by A earlier in the game. A game starts always implicitly in a_1. In a game $a_1b_1a_2b_2...a_nb_n$, A_i/B_i denote sets of vertices visited by A/B up to step $i \leq n$. Insufficiency of Game 3.1 for infinite graphs is shown in [5], so we only motivate changes and rules of Game 3.2.

- *Only B poisoning vertices he visits* is insufficient. For instance, on Yablo graph $(\omega, <)$, with no semikernel, A wins then choosing always a vertex past all played earlier by B.
- *B choosing only from out-neighbours of the last choice of A* would often allow A to win by following a ray. The graph $(\omega, <)$ with each edge $i < j$ subdivided into three edges (boldfaced vertices are from the orginal $(\omega, <)$) has no semikernel, but A wins choosing any out-neighbour of the last vertex

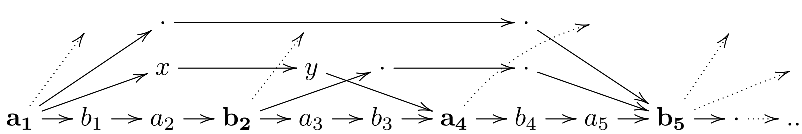

chosen by B. B must do the same, so each game on a dag follows a ray, here e.g., $a_1b_1a_2...$, and A wins if no edges connect ray's vertices except those of the ray itself (there may be paths, though). When B can return to earlier vertices, after a_4, which poisons y, he defeats A choosing x.
- *A can not choose freely* from out-neighbourhood of all B chosen vertices, since then A could avoid providing witness to some choices of B. The graph below has no semikernel, but with such a liberal rule A wins from a_1 choosing forever $1, 2, 3, ...$, also after B chose $\bullet$.

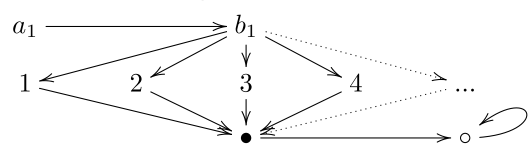

The game has a finitary character in the following sense. A wins if the game is infinite, while B when he chooses b_i with $\mathbf{E}(b_i) \subseteq \mathbf{E}^{\pm}(A_i) = \mathbf{E}(A_i) \cup \mathbf{E}^{-}(A_i)$, forcing A onto a poisoned vertex. This happens always after finitely many steps, when game traverses two paths from some (A visited) a to a_i and to a_j with $a_j \in \mathbf{E}(a_i)$ and A, having visited one of a_i, a_j, enters the other one.

A game is thus an ω sequence of vertices (moves) and is won by B if it contains a pair $a_j \in \mathbf{E}^{\pm}(a_i)$, for some $i < j < \omega$; otherwise it is won by A. Such games are ω games, and the former are treated as finite, terminating once A loses. We consider also transfinite games, but only to show that they are not needed. For an ordinal κ, a κ game is a game with κ moves of each player, where B starts after limit ordinals. A game is won by B if A visits a_i, a_j with $a_j \in \mathbf{E}(a_i)$, and by A if he never visits such a pair. (The ordinal notation is here particularly adequate, as well-foundedness reflects nonexistence of paths longer than ω. Each played vertex v is reached from a_1 by a finite path, found by going backwards from v in the ordinal sequence of moves.)

In general, the winner has no computable strategy but prescience – A knowing a semikernel (as will be shown) and B how to poison A in finite time (when graph is uncountable). We say that A/B *wins from* a start vertex a_1 if he can win using this prescience, no matter how the opponent plays. For A this means that he can win some game from a_1, while for B that there is no such game, i.e., that B wins every game which A starts from a_1.

Game's finitary character is crucial for our argument. The following two facts show that if B wins, on an arbitrary graph, he can do so in finitely many steps, hence for A it suffices to win an ω game.

FACT 3.3. *For each countable game, there is an ω game with the same moves.*

PROOF. More generally, for any limit ordinal λ, every $(\lambda+1)$ game, i.e., $...b_{\lambda+1}a_{\lambda+1}$, has an equivalent λ game (containing the same moves). In the fomer, B's last move is $a_i \to b_{\lambda+1}$ for some $i<\lambda$. Making this move just after a_i gives equivalent λ game $...a_ib_{\lambda+1}a_{\lambda+1}b_i...$ (where indices $_{\lambda+1}$ no longer reflect the sequence of moves). In the same way, $(\lambda+\omega)$ game can be turned into λ game by interleaving each $b_{\lambda+i}a_{\lambda+i}$ move, starting with $i=1$, from final ω into the preceding λ moves. Taking $\lambda=\omega$, for every countable game this gives an ω game with the same moves. □

Each countable transfinite game can be thus played with the same winner in ω steps. In particular, if B wins such a game, he can do it in finitely many steps. This holds also for uncountable graphs.

FACT 3.4. *On every sinkless graph, B wins iff he wins in finitely many steps.*

PROOF. For countable graphs the claim follows from Fact 3.3, so consider only uncountable graphs. If B wins in finitely many steps, then he obviously wins. The claim, and the proof, is that if B wins (every) uncountable game, he can also win (every game) in countably many steps. Hence, he wins in ω steps by Fact 3.3, and this happens only if A dies at some finite step.

So suppose that B wins (every) uncountable game from a_1, but loses some countable game g. This is due to B not having made some possible moves, so g can be extended by B making them. Let $C(g)$ be all such extensions of g to uncountable games, where B visits all $b\in\mathbf{E}(a)$, for each a visited by A. By assumption, B wins every $\hat{g}\in C(g)$, i.e., $\hat{g}$ contains a finite *failure* f, namely, two even paths $a_1b_{i1}...b_{in-1}a_{in}$ and $a_1b_{j1}...b_{jn-1}a_{jn}$, with $a_{in}\in\mathbf{E}(a_{jn})$. Adding these paths to g makes B win the resulting game, but we add only initial moves of $f\setminus g$, since A can now respond to them differently than in $\hat{g}$. These initial moves are B's, since in g each choice of B is already followed by a choice of A. So we let B make these initial moves from $f\setminus g$, A respond, and the game continue ω steps. By assumption, A wins countable game from a_1 no matter B's moves, hence A wins this extended game g_2. We extend it now in the same way as we extended g to g_2, choosing some finite failure f_2 of any uncountable extension $\hat{g}_2\in C(g_2)$, and obtaining a countable game g_3 won also by A. These extensions form an ordinal indexed chain of games $g\subset g_2\subset g_3\subset ...$, with A winning each game at countable index (because each step from g_i to g_{i+1} adds only countably many moves). In the limits λ, A wins game $g_\lambda=\bigcup_{i<\lambda}g_i$, as it is countable, but also because (*) if B wins it, he does so only through a finite failure, hence wins g_i for some $i<\lambda$. Eventually, the chain reaches the first uncountable ordinal ω_1, yielding uncountable game $g_{\omega_1}=\bigcup_{i<\omega_1}g_i$. A wins it by the argument (*), as ω_1 is a limit ordinal and A wins all games below it. But all uncountable games are won by B, so this contradiction establishes the claim. (Axiom of choice is needed to choose $\hat{g}_i$s and failures f_is.) □

By the next fact, A's victory is equivalent to exsitence of a semikernel. A minor technical difference, relevant later on, is that A survives Game 3.1 iff there is a semikernel, while Game 3.2 iff there is a semikernel containing the start vertex.

FACT 3.5. *On a sinkless G, A wins an ω game from a_1 iff G has a semikernel containing a_1.*

PROOF. If there is a semikernel $L\ni a_1$, A wins (ω game) choosing always a vertex from L. For the converse, assume first G to be countable. If A wins (no matter B's moves), let B exhaust systematically all out-neighbours of all A choices. More precisely, a possibly infinite game enumerates vertices played by A and assume an enumeration of out-neighbours of each among these. B starts by choosing $b_1\in\mathbf{E}(a_1)$ and continues with a fresh out-neighbour of each $a_j, j<i$ (skipping exhausted $\mathbf{E}(a_j)$), before he begins choosing out-neighbours of a new a_i, as indicated:

$a_1: b_1, b_2, b_4, b_7, ... \in\mathbf{E}(a_1)$
$a_2: b_3, b_5, b_8, ... \in\mathbf{E}(a_2)$
$a_3: b_6, b_9, ... \in\mathbf{E}(a_3)$
$a_4: b_{10}, ... \in\mathbf{E}(a_4)$

By the rules, each $a_{i+1} \in \mathbf{E}(b_i)$. If A wins, B has covered out-neighbourhood of every A_ω vertex, while for each vertex in B_ω, A has provided an out-neighbour, so $\mathbf{E}(A_\omega) \subseteq B_\omega \subseteq \mathbf{E}^-(A_\omega)$. Since A survived, no vertex in A_ω is poisoned, $A_\omega \cap (\mathbf{E}(A_\omega) \cup \mathbf{E}^-(A_\omega)) = \emptyset$, i.e., A_ω is independent. Thus A_ω is a semikernel containing a_1.

For uncountable graphs, by Fact 3.4 A wins iff he wins in ω steps. The game can then continue transfinitely to let B cover all $\mathbf{E}(a)$ (using axiom of choice) for every a chosen by A. Since A survives, this determines a semikernel as in the countable case. □

Two more lemmas are needed. The first uses an equivalent representation of a kernel K of G as a valuation of vertices $\kappa \in \mathbf{2}^{\mathbf{V}_G}$ given by (a) $\kappa(v) = \mathbf{1} \Leftrightarrow v \in K$ and such that (b) $\kappa(v) = \mathbf{1} \Leftrightarrow \forall y \in \mathbf{E}_G(v) : \kappa(y) = \mathbf{0}$. Conversely, each valuation satisfying (b) determines a kernel by (a). We now view kernels as such valuations. For a subgraph F of G, let $kerr(F)$ denote kernels of F relative to all valuations of its border, $\delta(F) = \mathbf{E}_G(\mathbf{V}_F) \setminus \mathbf{V}_F$. For any $X \subseteq \mathbf{2}^{\mathbf{V}_F \cup \delta(F)}$, let $X^* = X \times \mathbf{2}^{\mathbf{V}_G \setminus (\mathbf{V}_F \cup \delta(F))}$ denote all extensions of X to the whole G.

Lemma 3.6. *Let $\mathcal{H}$ be a set of induced subgraphs of a graph G covering it, $\mathbf{V}_G = \bigcup_{H \in \mathcal{H}} \mathbf{V}_H$, and such that for each finite $\mathcal{F} \subseteq \mathcal{H}$, the induced subgraph $G[\mathcal{F}] = G[\bigcup_{H \in \mathcal{F}} \mathbf{V}_H]$ is solvable. For each $x \in \mathbf{V}_G$, one of $\rho(x) = \mathbf{1}$ or $\rho(x) = \mathbf{0}$ can be then extended to $kerr^*(\mathcal{F})$, for each finite $\mathcal{F} \subseteq \mathcal{H}$:*
$\forall x \in \mathbf{V}_G \ \exists \rho \in \{\mathbf{1}, \mathbf{0}\}^x \ \forall \mathcal{F} \in \mathcal{P}^{fin}(\mathcal{H}) : kerr^*(G[\mathcal{F}]) \cap \{\rho\}^* \neq \emptyset.$

Proof. For every $x \in \mathbf{V}_G$ and finite $\mathcal{F} \subseteq \mathcal{H} : kerr^*(G[\mathcal{F}]) \cap (\{\mathbf{1}^x\}^* \cup \{\mathbf{0}^x\}^*) \neq \emptyset$. If for some $x \in \mathbf{V}_G$ and finite $\mathcal{A}, \mathcal{B} \subseteq \mathcal{H}$, (relative) kernels of $G[\mathcal{A}]$ and of $G[\mathcal{B}]$ can not agree on x, i.e., $kerr^*(G[\mathcal{A}]) \cap \{\mathbf{1}^x\}^* = \emptyset$ and $kerr^*(G[\mathcal{B}]) \cap \{\mathbf{0}^x\}^* = \emptyset$ (or dually), then

$$\begin{aligned} kerr^*(G[\mathcal{A} \cup \mathcal{B}]) &= kerr^*(G[\mathcal{A} \cup \mathcal{B}]) \cap (\{\mathbf{0}^x\}^* \cup \{\mathbf{1}^x\}^*) \\ &= kerr^*(G[\mathcal{A}]) \cap kerr^*(G[\mathcal{B}]) \cap (\{\mathbf{0}^x\}^* \cup \{\mathbf{1}^x\}^*) \\ &= (kerr^*(G[\mathcal{A}]) \cap kerr^*(G[\mathcal{B}]) \cap \{\mathbf{0}^x\}^*) \cup \\ &\quad\ (kerr^*(G[\mathcal{A}]) \cap kerr^*(G[\mathcal{B}]) \cap \{\mathbf{1}^x\}^*) = \emptyset \end{aligned}$$

contradicting solvability of $G[\mathcal{A} \cup \mathcal{B}]$. □

Conditions of the lemma allow us to find or produce a vertex v for which $\rho(v) = \mathbf{1}$ can be extended to $\overline{\rho} \in kerr^*(\mathcal{F})$, for every finite $\mathcal{F} \subseteq \mathcal{H}$. This can be one of vertices in any finite $\{v\} \cup \mathbf{E}_G(v)$. If $\mathbf{E}_G(v)$ is infinite for every $v \in \mathbf{V}_G$, we just add a new vertex x with an edge to arbitrary $y \in \mathbf{V}_G$. This does not change (non)existence of any semikernels, providing one $v \in \{x, y\}$ with $\rho(v) = \mathbf{1}$. Any such vertex is a *$\mathbf{1}$-vertex for $\mathcal{H}$*.

A sinkless graph can be covered by its sinkless induced subgraphs with finitely many ends.[2]

Lemma 3.7. *For each sinkless graph H there is a set $\mathcal{H}$ of its induced subgraphs such that*
(i) $\mathcal{H}$ covers H, i.e., $(\bigcup_{H_i \in \mathcal{H}} \mathbf{V}_{H_i}, \bigcup_{H_i \in \mathcal{H}} \mathbf{E}_{H_i}) = H$,
(ii) $\mathcal{H}$ is closed under finite unions: induced subgraph $H[\bigcup_1^n \mathbf{V}_{H_i}]$ is in $\mathcal{H}$, if $\{H_1, ..., H_n\} \subset \mathcal{H}$,
(iii) each $H_x \in \mathcal{H}$ is sinkless and has finitely many ends.

Proof. Every strictly $\preceq$-increasing sequence of ends has the supremum end by [15], Theorem 3.1. For each $\preceq$-maximal end E of H select a ray R_E with $\mathbf{V}_E = \mathbf{E}_H^*(R_E)$ and let $\mathcal{R}$ be the the set of such rays, so that $\mathbf{V}_H = \bigcup_{R \in \mathcal{R}} \mathbf{E}_H^*(R)$. For each $R_E \in \mathcal{R}$ and path π to R_E, let $V_E^\pi = \mathbf{V}_{R_E} \cup \mathbf{V}_\pi$ and VP be the set of all such, $VP = \{V_E^\pi \mid R_E \in \mathcal{R}, \pi \text{ is a path terminating in } R_E\} \subset \mathcal{P}(\mathbf{V}_H)$.

Set $\mathcal{H} = \{H[\bigcup_{V_E^\pi \in S} V_E^\pi] \mid S \in \mathcal{P}^{fin}(VP)\}$ satisfies then (i) and (ii) by construciton, and we show (iii). Each $H_x \in \mathcal{H}$ is sinkless, being a union of finitely many rays, each with finitely many paths to it. More precisely, there is a finite set J, such that $\mathbf{V}_{H_x} = \bigcup_{j \in J} (\mathbf{V}_{R_j} \cup P_j)$ where $R_j \in \mathcal{R}$, for each $j \in J$, and P_j is the set of vertices on finitely many paths to R_j. H_x has $|J|$ ends because each ray in H_x determines the same end of H_x as R_j for some $j \in J$, as we now argue.

Two *cofinal* rays, i.e., sharing a tail, determine the same end. Now, if finite union $H[\bigcup_{j \in J} H_j]$ of subgraphs from $\mathcal{H}$ contains an induced ray R that is not cofinal with any of the rays in any H_j, then $\mathbf{V}_R \cap \mathbf{V}_{H_i}$ is infinite for at least two $i \in J$. Consequently, no such ray is induced by

(a) $P = \bigcup_{j \in J} P_j$, since J and each P_j is finite,

[2]The characterization of set $\mathcal{H}$ in the following lemma corrects the original proposal. Its insufficiency for the proof of Theorem 3.8 was pointed out by Karl Heuer.

(b) P and any R_j, since P is finite,
(c) $\bigcup_{j \in J} \mathbf{V}_{R_j}$, since for each $i, k \in J$, rays R_i, R_k determine $\preceq$-incomparable ends, so all edges from R_i to the remaining rays $R_k, k \in J \setminus \{i\}$, start from a finite prefix of R_i, containing possible intersection with any ray induced by the union and not cofinal with R_i. □

As an example, consider the following graph, with edges $(a_i, a_{i+1}), (b_i, b_{i+1}), (b_i, a_{i+1})$ and (b_2, a_2).

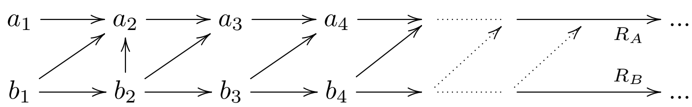

End A, determined by, e.g., ray $R_A = a_1 a_2 ...$, is here $\preceq$-maximal, so let $\mathcal{R} = \{R_A\}$. Finite subsets of $\{\{b_i, b_{i+1}, ..., b_{i+k}, a_{i+k+1}\} \mid i \geq 1, k \geq 0\} \cup \{\{b_1, b_2, a_2\}, \{b_2, a_2\}\}$, extended with $\mathbf{V}_{R_A}$, induce then members of $\mathcal{H}$. Ray $R_B = b_1 b_2 ...$ is contained in VP only through its finite approximations.

Theorem 3.8. *A graph is KP if each induced subgraph with finitely many ends is.*

Proof. Given agraph G, we remove its sinks as in the paragraph introducing sinks below Example 2.2, and assume G to be sinkless. Supposing that G is not KP, let H be its induced subgraph having no semikernel (existing by [7], Theorem 2). H is sinkless because each path starting from each $v \in \mathbf{V}_H$ extends to a ray. Otherwise some path in H would end with a rayless strong component C. If C has a kernel it is a semikernel of H, while unsolvable C is an induced subgraph of G with finitely many (zero) ends, contradicting KP-ness of all such.

Let $\mathcal{H}$ be a set of the induced subgraphs of H given by Lemma 3.7. Each $H_x \in \mathcal{H}$ is an induced subgraph of H, hence of G. Having finitely many ends, it is KP by the main assumption.

Let $a_1 \in \mathbf{V}_H$ be any $\mathbf{1}$-vertex for $\mathcal{H}$, so that $\rho(v) = \mathbf{1}$ can be extended to a (relative) kernel for each finite $\{H_1, ..., H_n\} \subset \mathcal{H}$. Let A start a game from a_1 and B choose b_1 in some $H_1 \in \mathcal{H}$. No matter how B plays, A wins the game if it stays inside H_1, so B must switch it to some other $H_1' \in \mathcal{H}$, i.e., to some vertex $b \in \mathbf{V}_{H_1'} \setminus \mathbf{V}_{H_1}$. The game now continues in the graph $(H_1 \cup H_1') = H_2 \in \mathcal{H}$. Again, if it stays there, A wins it no matter how B plays, so B must switch it to some vertex in H_2', outside of H_2. This holds for every finite number of steps, so that to win, B must visit infinitely many graphs from $\mathcal{H}$. But then the game is infinite and A wins it.

One might worry that earlier game on H_i, which A would win if it stayed there, may fail to extend monotonically to H_{i+1}, that is, although A wins from a_1 also in H_{i+1}, his earlier moves in H_i prevent now his victory in H_{i+1}.

This is not the case because A never chooses any a from which B wins (continuation of the actual game) in some $H_n \in \mathcal{H}$. This holds for a_1, so assume inductively that playing in H_i this holds for all moves up to a_j. Suppose that B playing now b_j forces A to choose some a_{j+1} from which B wins in some H_n. Then already step a_j was such, i.e., B wins from a_j in H_n, by playing b_j, contrary to the induction hypothesis. □

Induced subgraph of a safe graph is safe, while safe graph with finitely many ends is KP by [12], Theorem 3.19. Thus every induced subgraph with finitely many ends of a safe graph is KP, yielding

Corollary 3.9. *Every safe graph is KP.*

Since GNF theory without any liar or Yablo pattern is consistent, every paradox contains one of these patterns. The following, weaker compactness property may be sometimes easier to use.

Corollary 3.10. *A graph is KP if (subgraph induced by) each finite set of its ends is.*

Proof. If for each finite set $E \subseteq \vec{\vec{G}}$ subgraph H_E of G induced by E is KP, then each induced subgraph H' of G with finite $\vec{\vec{H'}}$ is KP, since such an H' is an induced subgraph of some H_E. The claim follows thus from Theorem 3.8. □

This does not imply Corollary 3.9 because (subgraph induced by) only a single end may contain infinitely many other ends, for instance, a grid of infinitely many horizontal and vertical rays/ends.

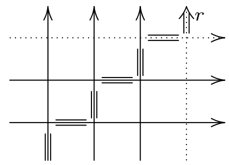

The end $\mathbf{E}^*_{=}(r)$, for any 'diagonal' ray r, comprises the whole graph and all other ends. Theorem 3.19 from [12] does not apply to such a graph: even if safe and induced by a single end, it possesses infinitely many ends. Corollary 3.10 trivializes for it, while Theorem 3.8 gives its solvability.

The notion of an induced subgraph translates naturally into language of logic as removal of atoms. For a GNF theory over atoms $\mathbf{V}$, a *subtheory induced by* $X \subseteq \mathbf{V}$ results by removing axiom $y \leftrightarrow ...$, for each $y \in \mathbf{V} \setminus X$, and then removing atoms $\mathbf{V} \setminus X$ from (the right sides of) all remaining axioms. KP-ness amounts then to a stronger form of consistency, say, *hereditary consistency*, HC, holding when every induced subtheory is consistent. An example of a consistent theory which is not HC is contingent liar $\{x \leftrightarrow \neg x \wedge \neg y, y \leftrightarrow \neg \overline{y}, \overline{y} \leftrightarrow \neg y\}$, with the graph $\circlearrowleft x \to y \leftrightarrows \overline{y}$ and inconsistent subtheory $\{x \leftrightarrow \neg x\}$ induced by $\{x\}$.

Translating rays and ends by mere copying graph definitions – a ray as an infinite chain of axiomatic implications (from left to right of (2.1)) $x_i \to \neg x_{i+1}, i \in \omega$, with $x_i \neq x_j$ for $i \neq j$, and an end as all such finite chains of implications to a ray – the two compactness theorems say then that a GNF theory Γ is HC if

3.8: each induced subtheory of Γ with finitely many ends is HC, or

3.10: subtheory induced by each finite set of ends of Γ is HC.

Acknowledgments. Thanks to Sjur Dyrkolbotn for suggesting improvements to a draft of the proof, to Karl Heuer for pointing out a mistake there, and to L. Meltzers Høyskolefond for financial support.

REFERENCES

[1] T. Beringer and T. Schindler. A graph-theoretic analysis of the semantic paradoxes. *The Bulletin of Symbolic Logic*, 23(4):442–492, 2017.

[2] M. Bezem, C. Grabmayer, and M. Walicki. Expressive power of digraph solvability. *Annals of Pure and Applied Logic*, 163(3):200–212, 2012.

[3] T. Bolander. *Logical theories for agent introspection*. PhD thesis, IMM, Technical University of Denmark, 2003.

[4] R. Cook. Patterns of paradox. *The Journal of Symbolic Logic*, 69(3):767–774, 2004.

[5] P. Duchet and H. Meyniel. Kernels in directed graphs: a poison game. *Discrete Mathematics*, 115(1-3):273–276, 1993.

[6] H. Gaifman. Pointers to truth. *The Journal of Philosophy*, 89(5):223–261, 1992.

[7] V. Neumann-Lara. Seminúcleos de una digráfica. Technical report, Anales del Instituto de Matemáticas II, Universidad Nacional Autónoma México, 1971.

[8] L. Rabern, B. Rabern, and M. Macauley. Dangerous reference graphs and semantic paradoxes. *Journal of Philosophical Logic*, 42:727–765, 2013.

[9] M. Richardson. Solutions of irreflexive relations. *The Annals of Mathematics, Second Series*, 58(3):573–590, 1953.

[10] R. Sorensen. Yablo's paradox and kindred infinite liars. *Mind*, 107:137–155, 1998.

[11] M. Walicki. Resolving infinitary paradoxes. *Journal of Symbolic Logic*, 82(2):709–723, 2017.

[12] M. Walicki. Kernels of digraphs with finitely many ends. *Discrete Mathematics*, 342:473–486, 2019.

[13] M. Walicki. Extensions in graph normal form. *Logic Journal of the IGPL*, 30(1):101–123, 2022.

[14] S. Yablo. Circularity and paradox. In T. Bolander, V. E. Hendricks, and S. A. Pedersen, editors, *Self-Reference*. CSLI Publications, 2004.

[15] J. Zuther. Ends in digraphs. *Discrete Mathematics*, 184(1-3):225–244, 1998.

DEPARTMENT OF INFORMATICS
UNIVERSITY OF BERGEN
PBOX 7803, 502 BERGEN, NORWAY
E-mail: Michal.Walicki@uib.no