18739A: Foundations of Security and Privacy

Contextual Integrity and its Formalization

Anupam Datta
CMU

Fall 2007-08

Problem Statement

- Is an organization's business process compliant with privacy regulations and internal policies?

- Examples of organizations
  - Hospitals, financial institutions, other enterprises handling sensitive information
- Examples of privacy regulations
  - HIPAA, GLBA, COPPA, SB1386

Goal: Develop methods and tools to answer this question

Contextual Integrity [N2004]

- Philosophical framework for privacy
- Central concept: *Context*
  - Examples: Healthcare, banking, education
- What is a context?
  - Set of interacting *agents* in *roles*
    - Roles in healthcare: doctor, patient, …
  - *Norms of transmission*
    - Doctors should share patient health information as per the HIPAA rules
  - *Purpose*
    - Improve health

Outline

1. Motivating Example 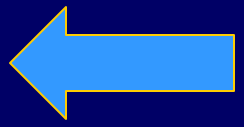
2. Framework
   - Model
   - Logic of Privacy and Utility
3. Workflows and Responsibility
4. Algorithmic Results
   - Workflow Design assuming agents responsible
   - Auditing logs when agents irresponsible
5. Conclusions

MyHealth@Vanderbilt Workflow

Humans + Electronic system

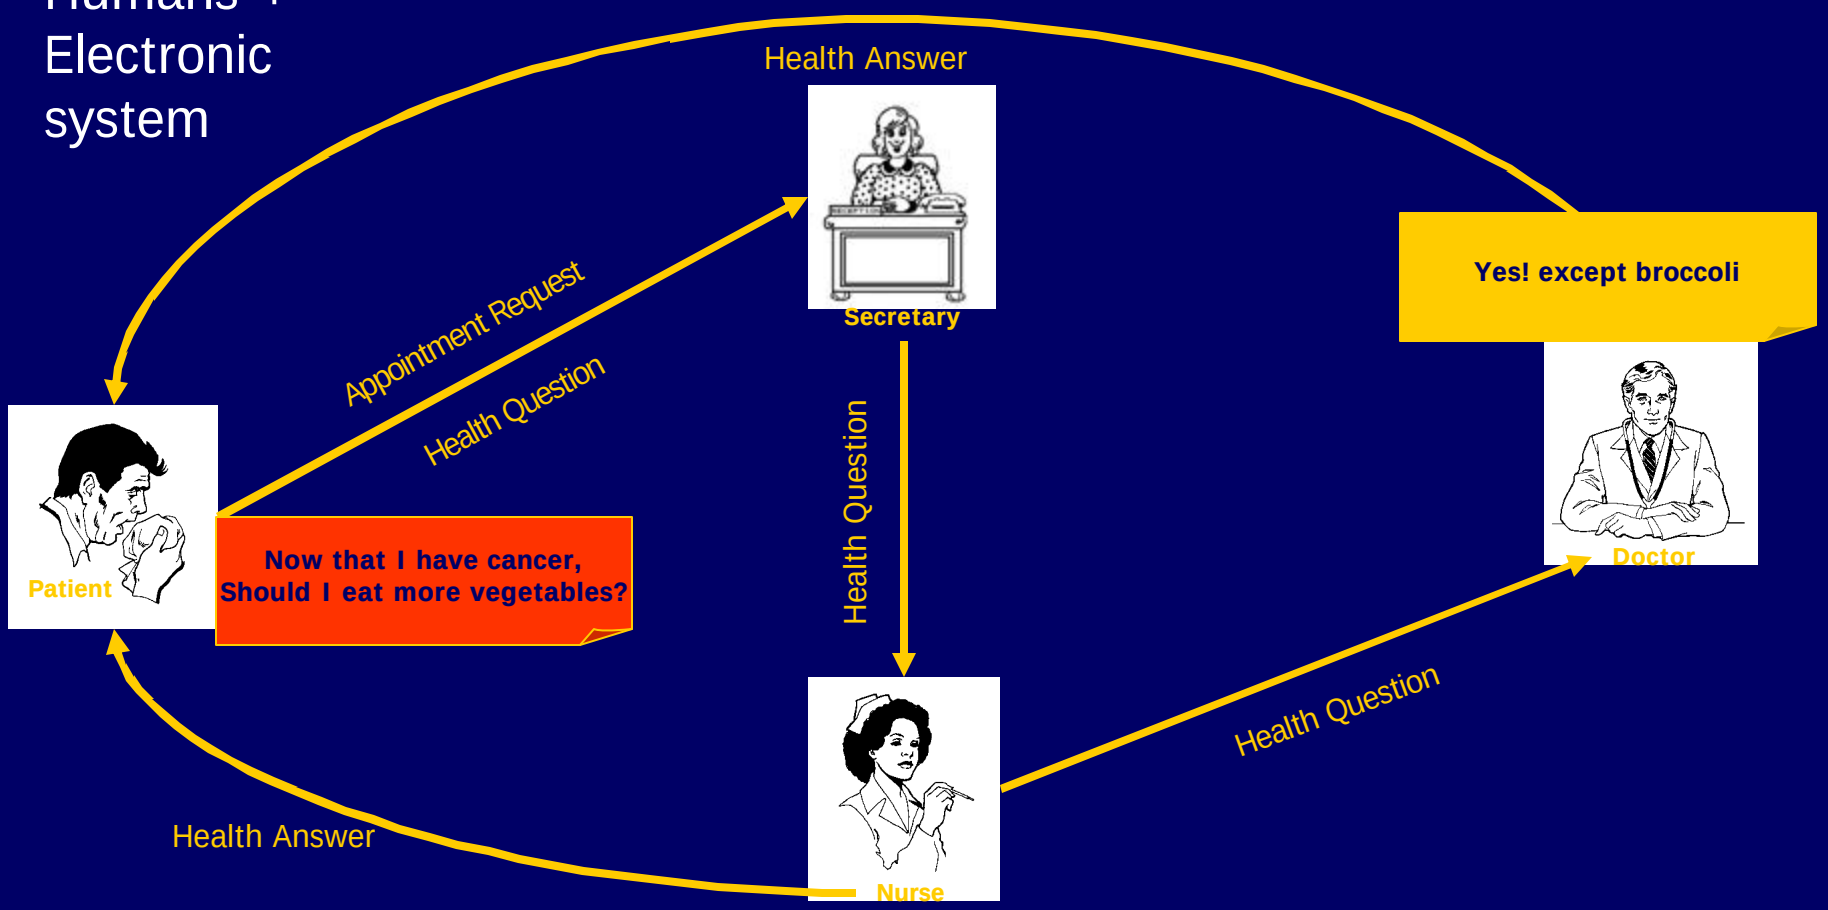

Possible Enhancement

- Secretary handles every message
  - Privacy
  - Efficiency
  - Robustness
- Messages opaque to MyHealth
  - Unable to help secretary route messages
  - Hinders adding features like delegation
- Suggestion: add short tags to messages

MyHealth with Message Tags

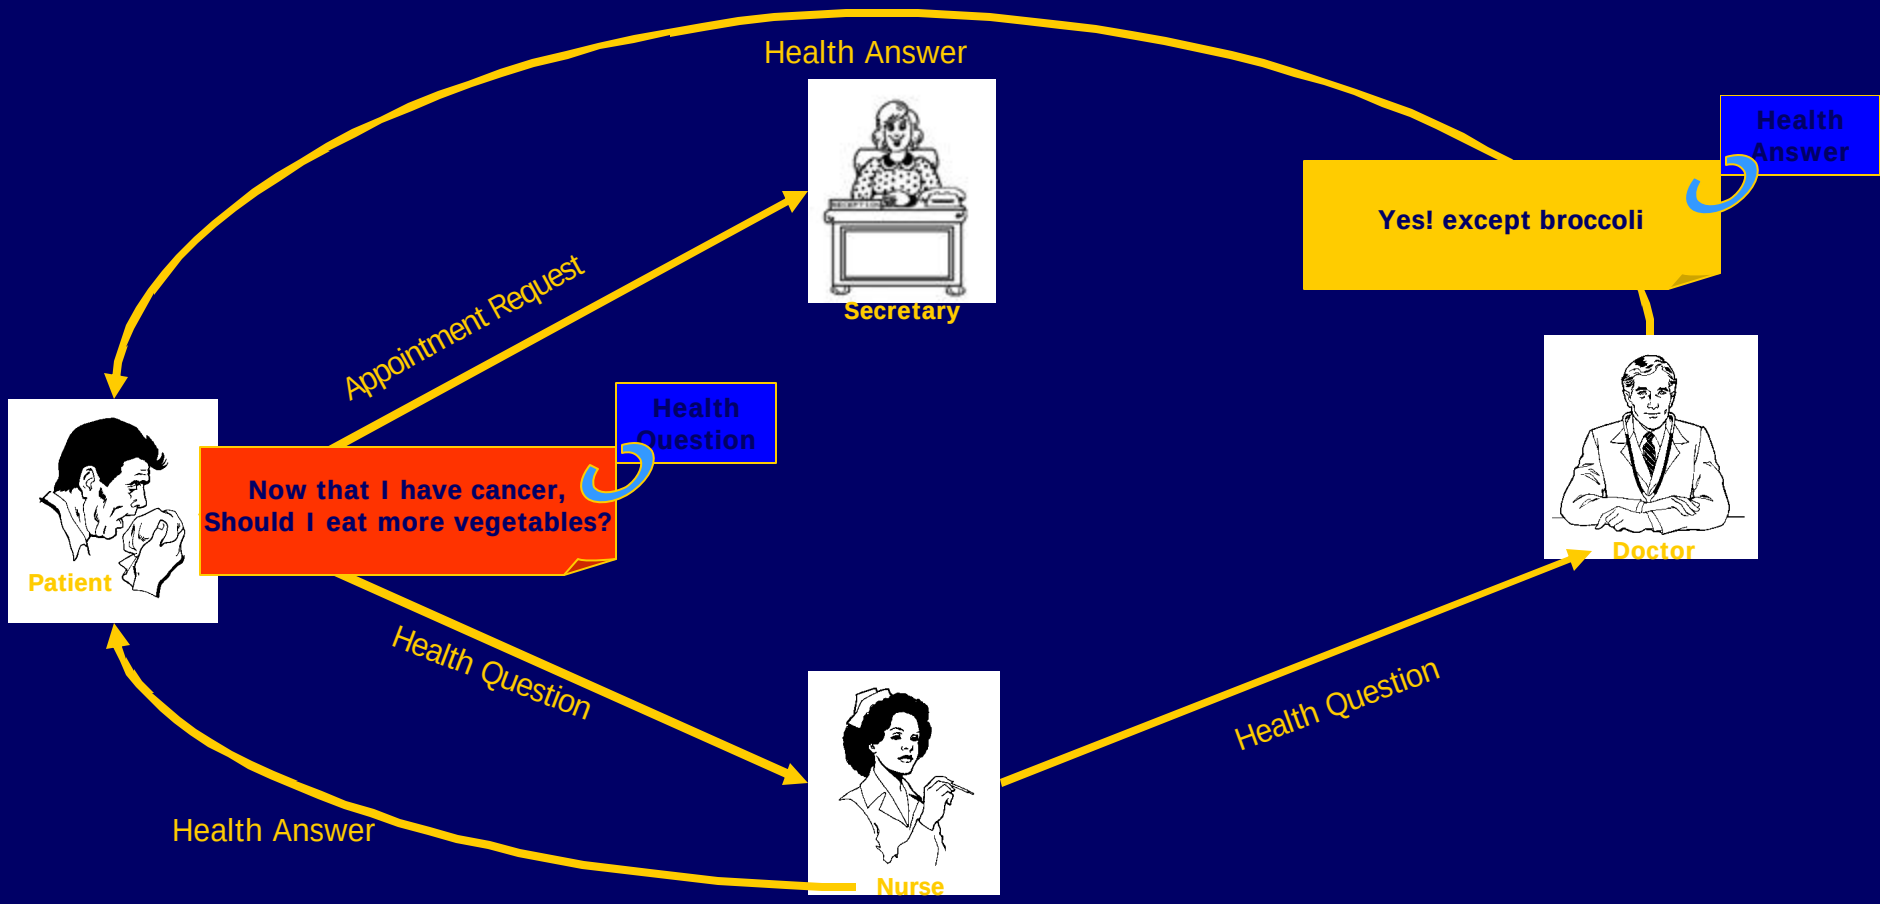

Robustness to Mistaken Tags

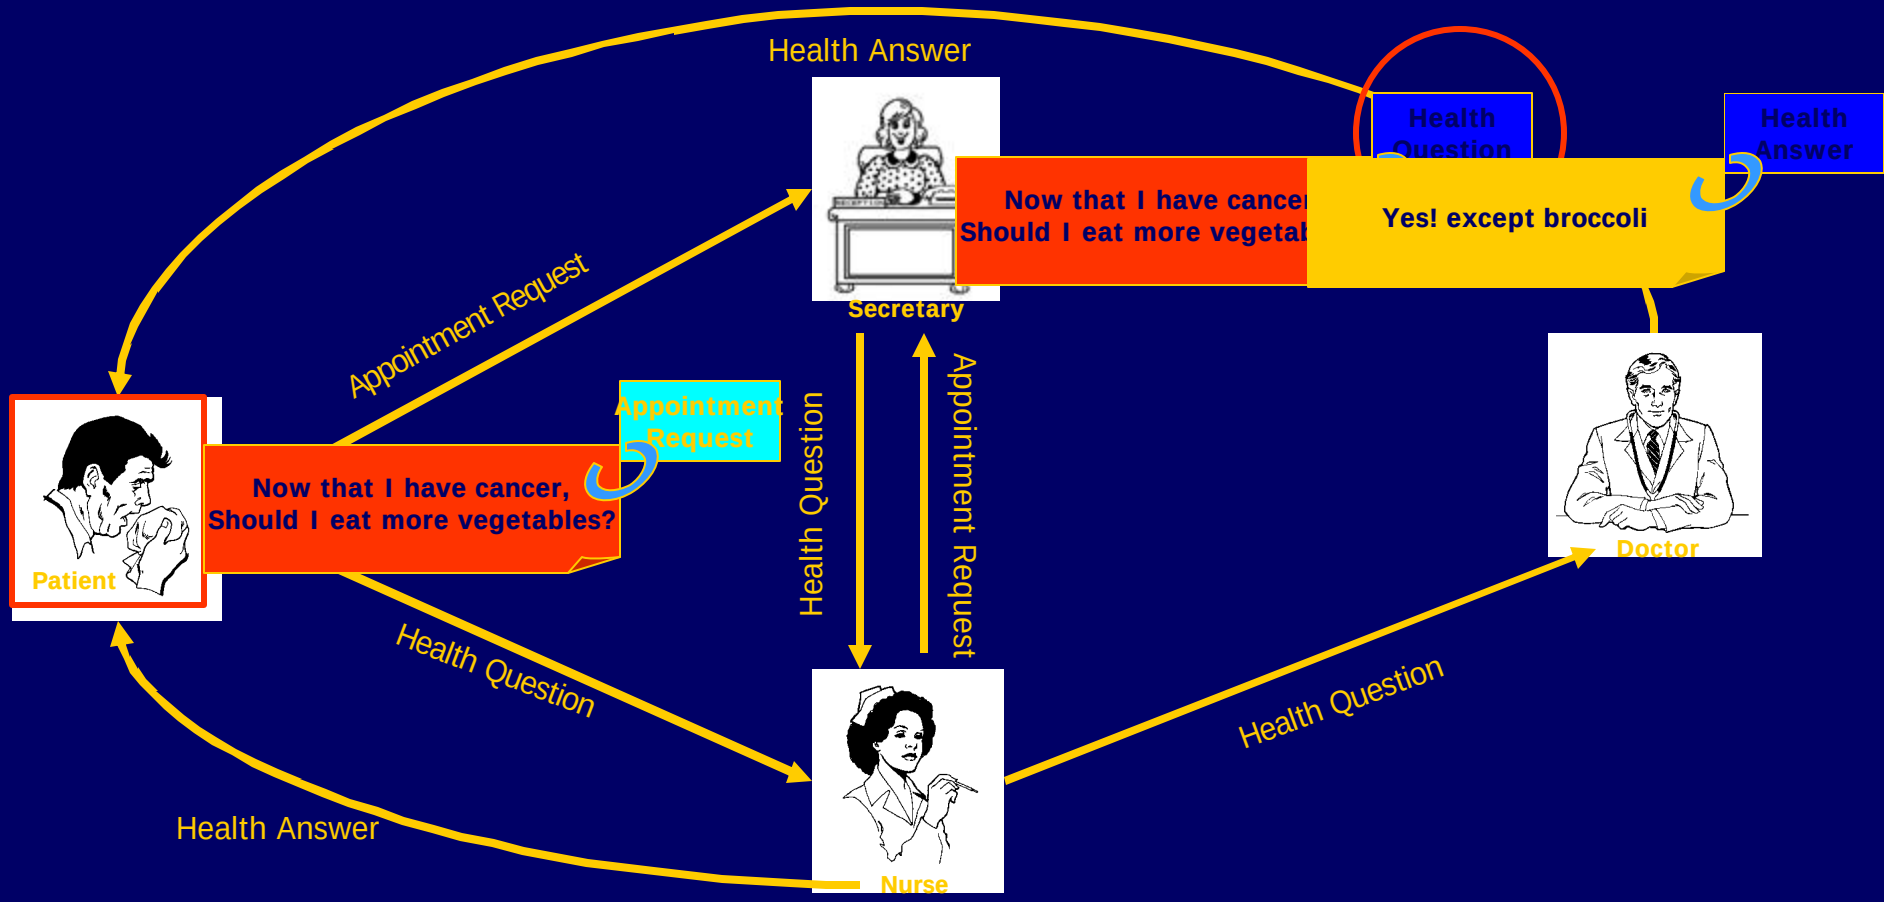

Workflow Design Goals

- Privacy
  - Secretary does not get sensitive info
- Utility
  - Health question eventually answered
- Robustness
  - Properties hold even with mistakes

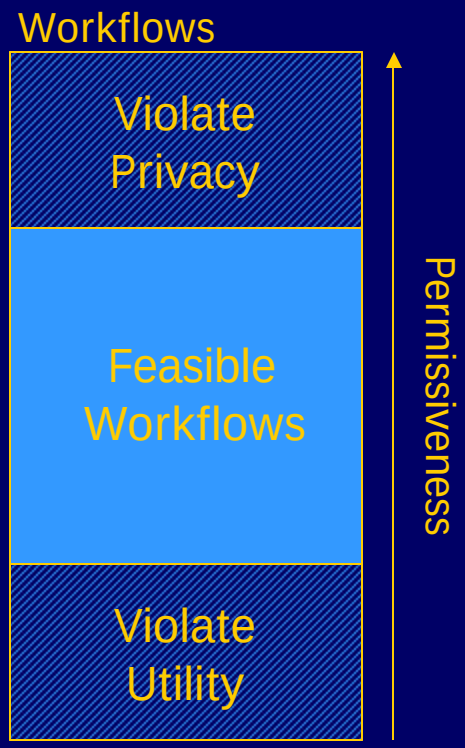

Recommendations

- Add short tags to messages
  - Enhances privacy
  - Increases efficiency
  - Scales with added functionality
- Assign responsibilities
  - Example: secretary should tag messages with “health question” if needed

Outline

1. Motivating Example
2. Framework 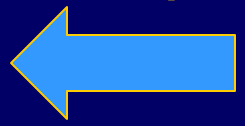
   - Model
   - Logic of Privacy and Utility
3. Workflows and Responsibility
4. Algorithmic Results
   - Workflow Design assuming agents responsible
   - Auditing logs when agents irresponsible
5. Conclusions

Informational Norms

"In a context, the flow of information of a certain type about a subject (acting in a particular capacity/role) from one actor (could be the subject) to another actor (in a particular capacity/role) is governed by a particular transmission principle."

Contextual Integrity [N2004]

Model

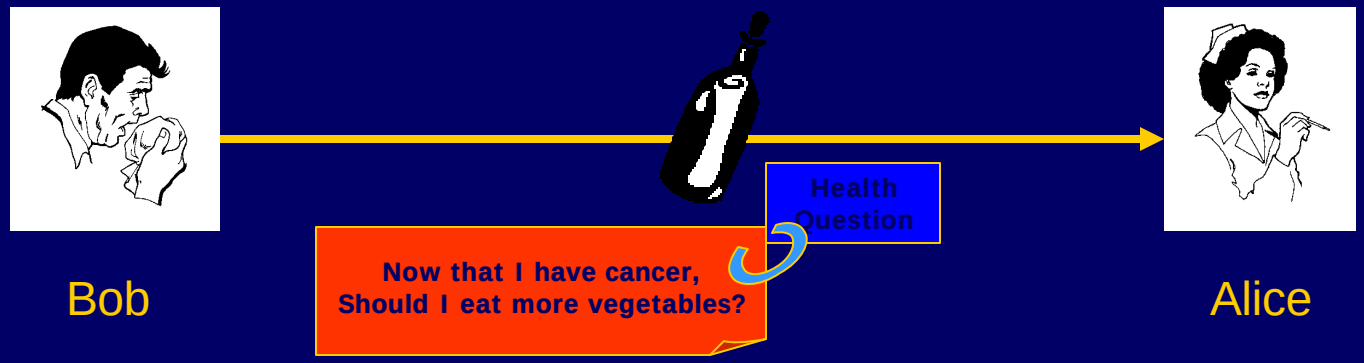

- Communication via send actions:
  - Sender: Bob in role Patient
  - Recipient: Alice in role Nurse
  - Subject of message: Bob
  - Tag: Health Question
  - Message: Now that ….
- Data model & knowledge evolution:
  - Agents acquire knowledge by:
    - receiving messages
    - deriving additional attributes based on data model
      - Health Question ≤ Protected Health Information

contents(msg) vs. tags (msg)

Model

[BDMS2007]

- *State* determined by knowledge of each agent
- *Transitions* change state
  - Set of concurrent send actions
  - Send(p,q,m) possible only if agent p knows m

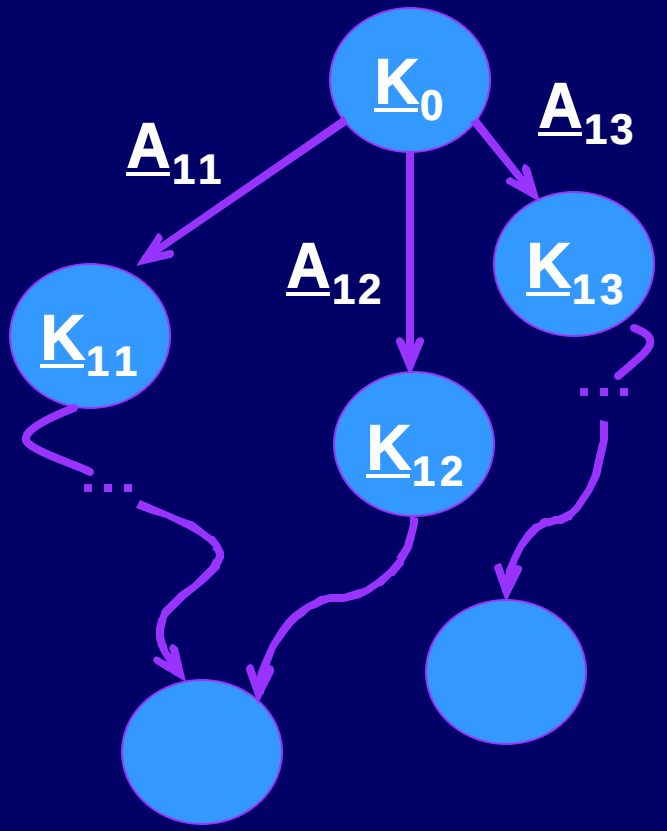

Concurrent Game Structure

G = <k, Q, Π, π, d, δ>

Logic of Privacy and Utility

- Syntax

j ::= send(p_1,p_2,m)	p_1 sends p_2 message m
| contains(m, q, t)	m contains attrib t about q
| tagged(m, q, t)	m tagged attrib t about q
| inrole(p, r)	p is active in role r
| $t \leq t'$	Attrib t is part of attrib t'
| $j \wedge j$ | $\neg j$ | $\exists x.\ j$	Classical operators
| $j \mathrm{U} j$ | $j \mathrm{S} j$ | $\mathrm{O} j$	Temporal operators
| $\langle\langle \underline{p} \rangle\rangle j$	Strategy quantifier

- Semantics

  Formulas interpreted over concurrent game structure

Specifying Privacy

- MyHealth@Vanderbilt

  In all states, only nurses and doctors receive health questions

$$G \; \forall \; p1, p2, q, m$$
$$send(p1, p2, m) \wedge contains(m, q, health\text{-}question)$$
$$\Rightarrow inrole(p2, nurse) \vee inrole(p2, doctor)$$

LTL fragment can express HIPAA, GLBA, COPPA [BDMN2006]

Specifying Utility

- MyHealth@Vanderbilt

Patients have a strategy to get their health questions answered

$$\forall\, p\ \text{inrole}(p, \text{patient}) \Rightarrow$$
$$\langle\langle p \rangle\rangle\, F\, \exists\, q, m.$$
$$\text{send}(q, p, m) \wedge \text{contains}(m, p, \text{health-answer})$$

Expressing Privacy

$$\Box\forall p_1, p_2, q : P.\forall m : M.\forall t : T.$$

$$\text{incontext}(p_1, c) \wedge \text{send}(p_1, p_2, m) \wedge \text{contains}(m, q, t) \rightarrow \bigvee_{\varphi^+ \in \text{norms}^+(c)} \varphi^+ \wedge \bigwedge_{\varphi^- \in \text{norms}^-(c)} \varphi^-$$

$$\text{positive norm:} \quad \text{inrole}(p_1, \hat{r}_1) \wedge \text{inrole}(p_2, \hat{r}_2) \wedge \text{inrole}(q, \hat{r}) \wedge (t \in \hat{t}) \wedge \theta \wedge \psi$$

$$\text{negative norm:} \quad \text{inrole}(p_1, \hat{r}_1) \wedge \text{inrole}(p_2, \hat{r}_2) \wedge \text{inrole}(q, \hat{r}) \wedge (t \in \hat{t}) \wedge \theta \rightarrow \psi$$

Figure 1. Norms of Transmission Represented as a Temporal Formula

Allow message transmission if:

- at least one positive norm is satisfied; and
- all negative norms are satisfied

HIPAA – Healthcare Privacy

$$\text{inrole}(p_1, \textit{covered-entity}) \wedge \text{inrole}(p_2, \textit{individual}) \wedge (q = p_2) \wedge (t \in \textit{phi}) \quad (2)$$

$$\text{inrole}(p_1, \textit{covered-entity}) \wedge \text{inrole}(p_2, \textit{provider}) \wedge \text{inrole}(q, \textit{patient}) \wedge (t \in \textit{phi}) \quad (3)$$

$$\text{inrole}(p_1, \textit{covered-entity}) \wedge \text{inrole}(p_2, \textit{individual}) \wedge (q = p_2) \wedge (t \in \textit{psychotherapy-notes}) \rightarrow$$
$$\diamondminus \exists p : P.\, \text{inrole}(p, \textit{psychiatrist}) \wedge \text{send}(p, p_1, \textit{approve-disclose-psychotherapy-notes}) \quad (4)$$

$$\text{inrole}(p_1, \textit{covered-entity}) \wedge \text{inrole}(p_2, \textit{individual}) \wedge \text{inrole}(q, \textit{individual}) \wedge (t \in \textit{condition-and-location}) \wedge$$
$$\diamondminus \exists m' : M.\, \text{send}(p_2, p_1, m') \wedge \text{contains}(m', q, \textit{name}) \quad (5)$$

$$\text{inrole}(p_1, \textit{covered-entity}) \wedge \text{inrole}(p_2, \textit{clergy}) \wedge \text{inrole}(q, \textit{individual}) \wedge (t \in \textit{directory-information}) \quad (6)$$

Figure 2. Norms of Transmission from the HIPAA Privacy Rule

- HIPAA consists primarily of positive norms: share phi if some rule explicitly allows it (2), (3), (5), (6)
- Exception: negative norm about psychotherapy notes (4)

COPPA – Children Online Privacy

$$\text{inrole}(p_1, \textit{child}) \wedge \text{inrole}(p_2, \textit{web-site}) \wedge (q = p_1) \wedge (t \in \textit{protected-info}) \rightarrow$$
$$\exists p : P.\, \text{inrole}(p, \textit{parent}) \wedge \neg \text{send}(p, p_2, \textit{revoke-consent}) \mathcal{S}$$
$$(\text{send}(p, p_2, \textit{grant-consent}) \wedge \diamondsuit \text{send}(p_2, p, \textit{privacy-notice})) \quad (7)$$

$$\text{inrole}(p_1, \textit{child}) \wedge \text{inrole}(p_2, \textit{web-site}) \wedge (q = p_1) \wedge (t \in \textit{protected-info}) \rightarrow$$
$$\Box \forall p : P.\, \text{inrole}(p, \textit{parent}) \wedge \text{send}(p, p_2, \textit{request-information}) \rightarrow$$
$$\diamondsuit (\text{send}(p_2, p, \textit{privacy-notice}) \wedge \text{send}(p_2, p, m)) \quad (8)$$

Figure 3. Norms of Transmission from COPPA

- COPPA consists primarily of negative norms
  - children can share their protected info *only if* parents consent (7) (condition)
  - (8) (obligation – future requirements)

GLBA – Financial Institutions

Sender role

Attribute

Subject role

Financial institutions must notify consumers if they share their non-public personal information with non-affiliated companies, but the notification may occur either before or after the information sharing occurs

Recipient role

Transmission principle

$$\text{inrole}(p_1, \textit{institution}) \wedge \text{inrole}(p_2, \textit{non-affiliate}) \wedge \text{inrole}(q, \textit{consumer}) \wedge (t \in \textit{npi}) \rightarrow$$
$$\Diamond \text{send}(p_1, q, \textit{privacy-notice}) \vee \Diamond\!\!-\text{send}(p_1, q, \textit{privacy-notice})$$

Outline

1. Motivating Example
2. Framework
   - Model
   - Logic of Privacy and Utility
3. Workflows and Responsibility 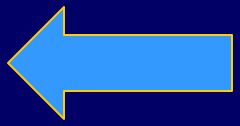
4. Algorithmic Results
   - Workflow Design assuming agents responsible
   - Auditing logs when agents irresponsible
5. Conclusions

MyHealth@Vanderbilt Improved

Doctor should answer health questions

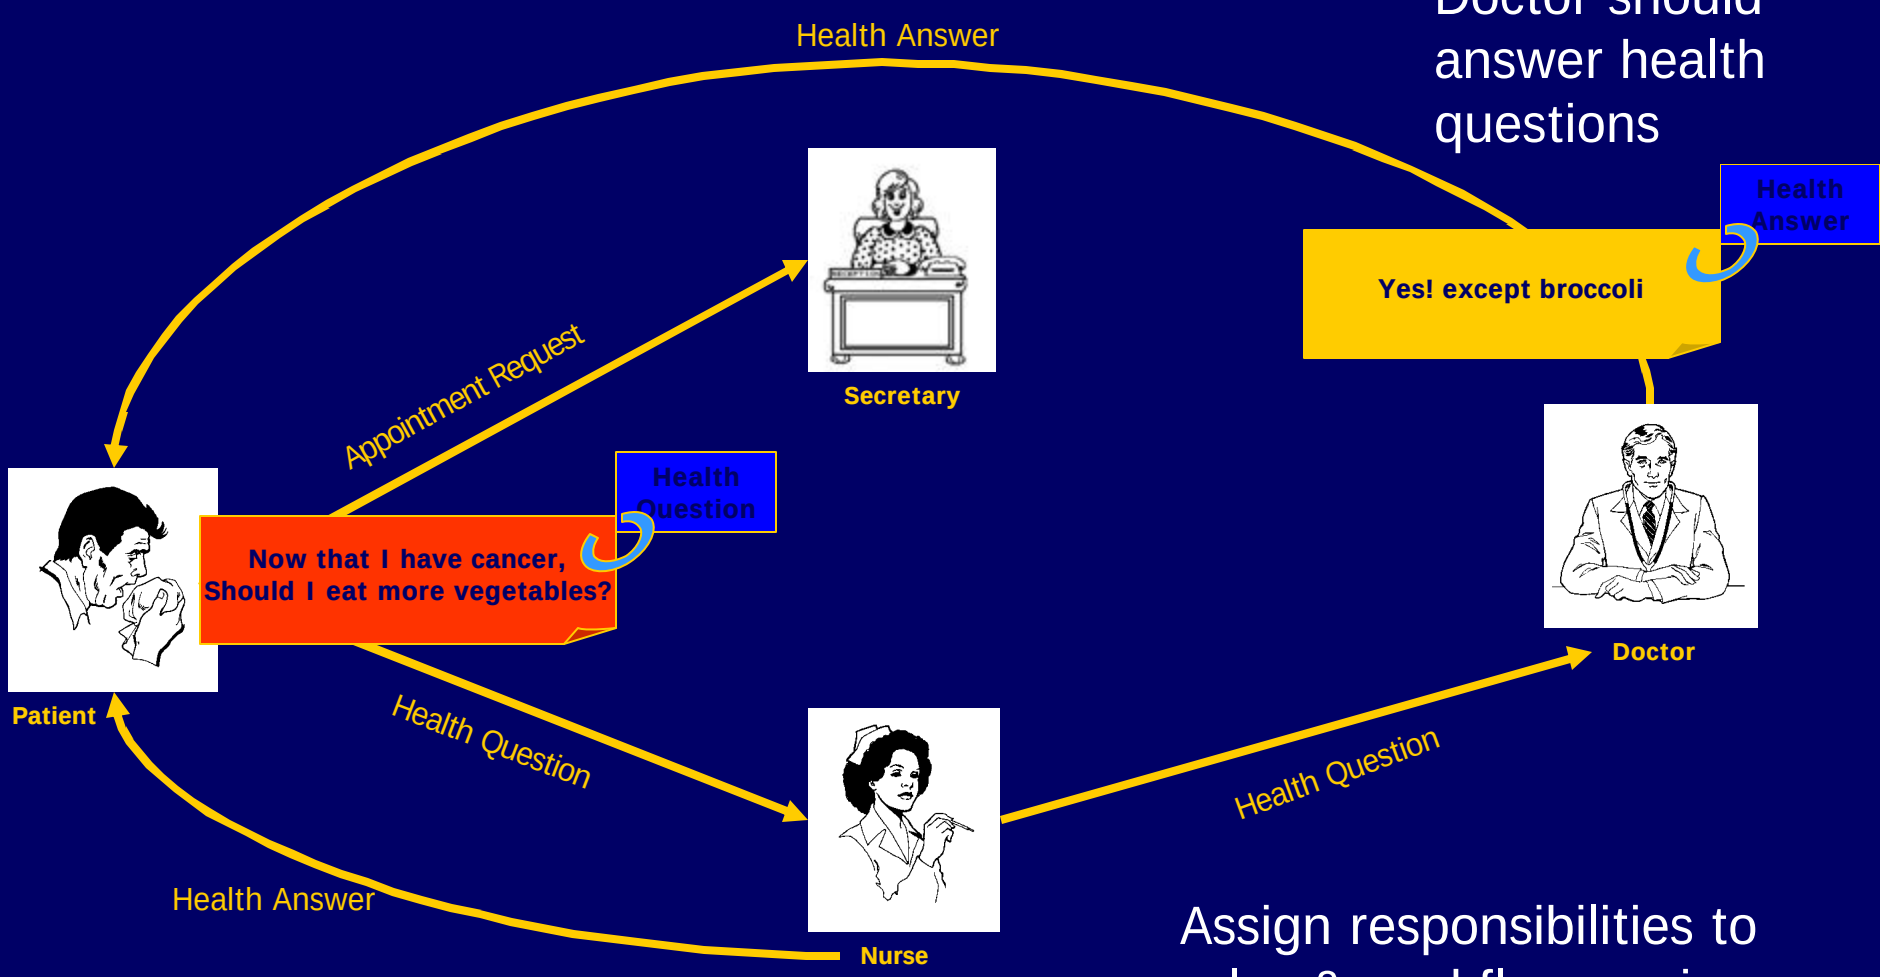

Assign responsibilities to roles & workflow engine

Graph-based Workflow

- Graph
  - (R, R x R), where R is the set of roles
- Edge-labeling function
  - permit: R x R $\rightarrow 2^T$, where T is the set of attributes
- Responsibility of workflow engine
  - Allow msg from role r_1 to role r_2 iff
  - tags(msg) $\subseteq$ permit(r_1, r_2)
- Responsibility of human agents in roles
  - Tagging responsibilities
    - ensure messages are correctly tagged
  - Progress responsibilities
    - ensure messages proceed through workflow

MyHealth Responsibilities

- Tagging
  - Nurses should tag health questions

    $G\ \forall p, q, s, m.\ \text{inrole}(p, \text{nurse}) \wedge \text{send}(p, q, m) \wedge \text{contains}(m, s, \text{health-question}) \Rightarrow \text{tagged}(m, s, \text{health-question})$
- Progress
  - Doctors should answer health questions

    $G\ \forall p, q, s, m.\ \text{inrole}(p, \text{doctor}) \wedge \text{send}(q, p, m) \wedge \text{contains}(m, s, \text{health-question}) \Rightarrow F\ \exists m'.\ \text{send}(p, s, m') \wedge \text{contains}(m', s, \text{health-answer})$

Abstract Workflow

- Responsibility of workflow engine
  - LTL formula j
  - *Feasible* (enforceable) if j is a safety formula without the contains() predicate

- Responsibility of each role r
  - LTL formula j_r
  - *Feasible* if agents have a strategy to discharge their responsibilities

$$\forall \text{ p. } j \wedge \text{inrole(p, r)} \Rightarrow \langle\langle \text{p} \rangle\rangle j_r$$

Graph-based workflows are a special case

Outline

1. Motivating Example
2. Framework
   - Model
   - Logic of Privacy and Utility
3. Workflows and Responsibility
4. Algorithmic Results
   - Workflow Design assuming agents responsible
     - Abstract workflows
   - Auditing logs when agents irresponsible
     - Only graph-based workflows
5. Conclusions

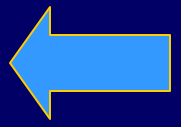

MyHealth@Vanderbilt Improved

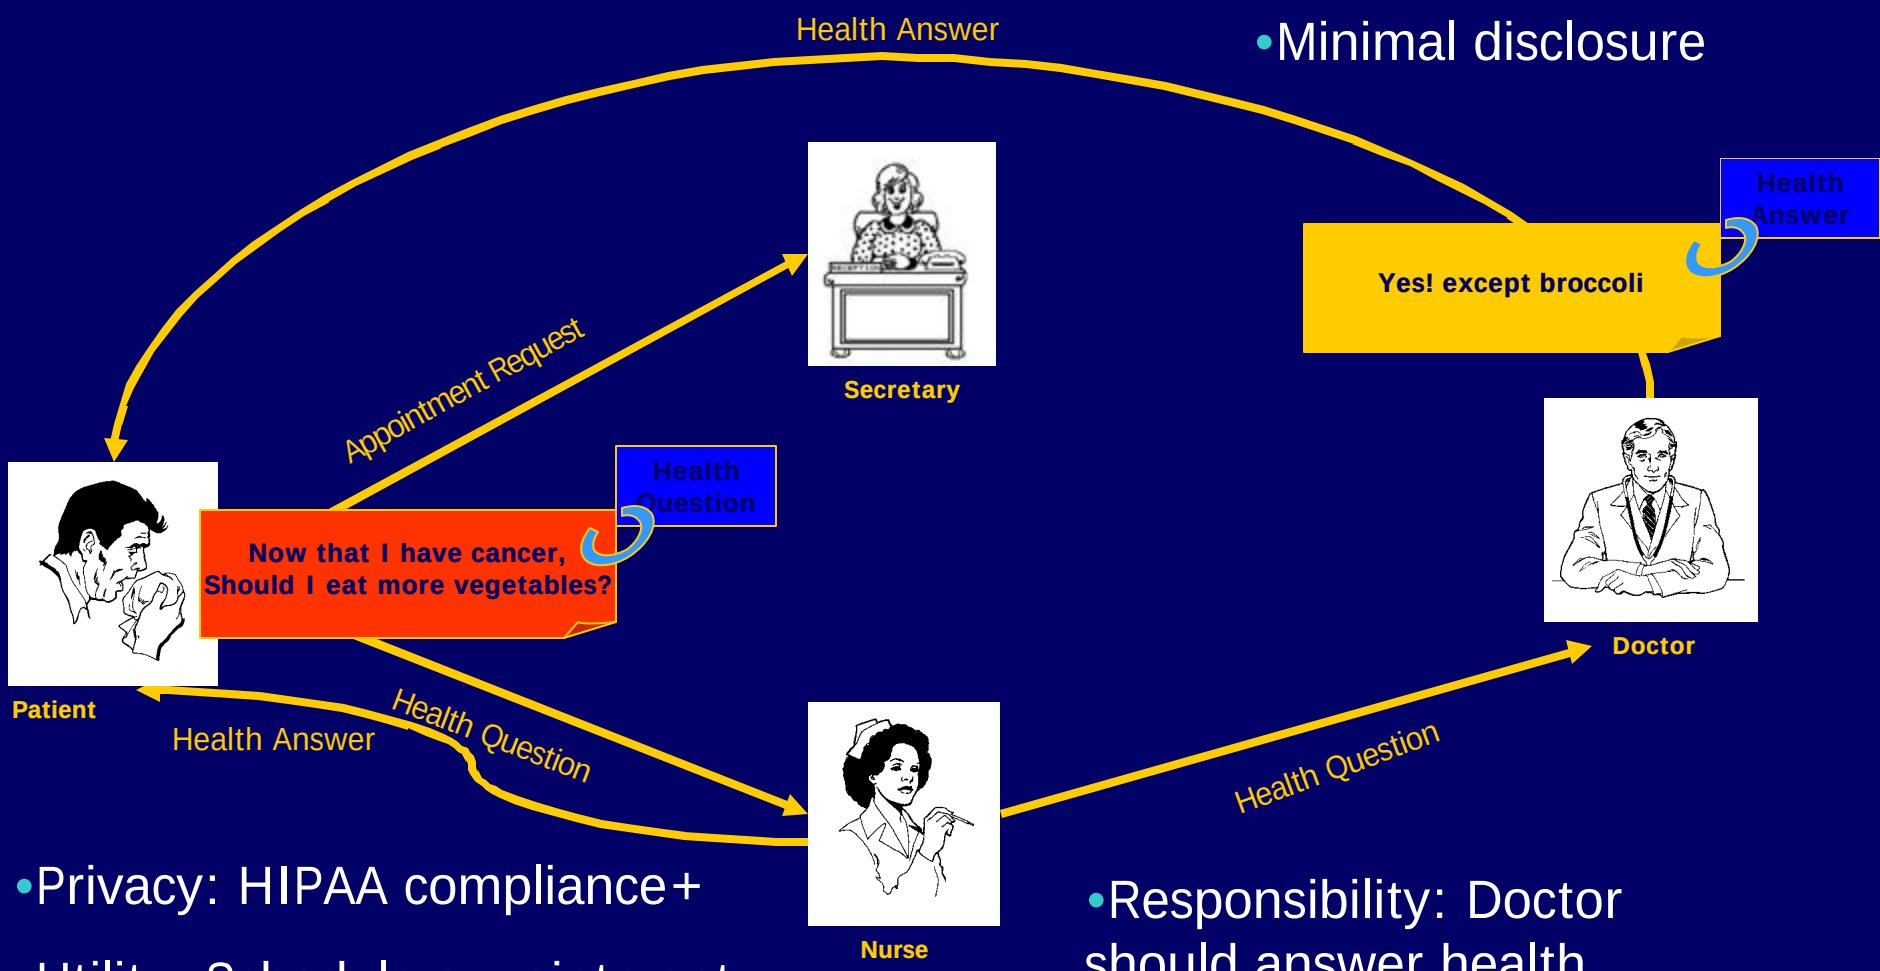

- Minimal disclosure
- Privacy: HIPAA compliance+
- Utility: Schedule appointments, obtain health answers
- Responsibility: Doctor should answer health questions

Workflow Design Results

- Theorems:

  Assuming all agents act responsibly, checking whether workflow achieves

  - Privacy is in PSPACE (in the size of the formula describing the workflow)
  - Utility is decidable

- Definition and construction of minimal disclosure workflow

Algorithms implemented in model-checkers, e.g. SPIN, MOCHA

Deciding Privacy

- PLTL model-checking problem is PSPACE decidable

  G |= tags-correct ***U*** agents-responsible
  $\Rightarrow$ privacy-policy

  G: concurrent game structure

Result applies to finite models (#agents, msgs,...)

MyHealth Privacy

- MyHealth@Vanderbilt workflow satisfies this privacy condition

  In all states, only nurses and doctors receive health questions

  $G \forall p1, p2, q, m$

  $send(p1, p2, m) \wedge contains(m, q, health\text{-}question)$

  $\Rightarrow inrole(p2, nurse) \vee inrole(p2, doctor)$

- Run LTL model-checker, e.g. SPIN

Deciding Utility

- ATL* model-checking of concurrent game structures is
  - Decidable with perfect information
  - Undecidable with imperfect information
- Theorem:

  There is a sound decision procedure for deciding whether workflow achieves utility
- Intuition:
  - Translate imperfect information into perfect information by considering possible actions from one player's point of view

MyHealth Utility

- MyHealth@Vanderbilt workflow satisfies this utility condition

  Patients have a strategy to get their health questions answered

  $\forall$ p inrole(p, patient) $\Rightarrow$

  <<p>> F $\exists$ q, m.

  send(q, p, m) $\wedge$ contains(m, p, health-answer)

- Run ATL* model-checker, e.g. MOCHA

Minimal Disclosure Workflow

- Abstract workflows:

$$W_1(\varphi_1, \varphi_R) \leq W_2(\varphi_2, \varphi_R) \text{ if}$$

$$\boldsymbol{G} \text{ satisfies } \varphi_1 \Rightarrow \varphi_2$$

- Graph-based workflows:

$$W_1(R, permit_1) \leq W_2(R, permit_2) \text{ if}$$

$$\forall\, r_1, r_2 \in R.\ permit_1(r_1, r_2) \subseteq permit_2(r_1, r_2)$$

- Lemma:

  If $W_1 \leq W_2$ and W_2 achieves a privacy goal, then so does W_1

- Minimal Disclosure Workflow:

  W is minimal wrt to a utility goal if W achieves the goal and all feasible W' < W fails to achieve the goal

MyHealth@Vanderbilt Workflow

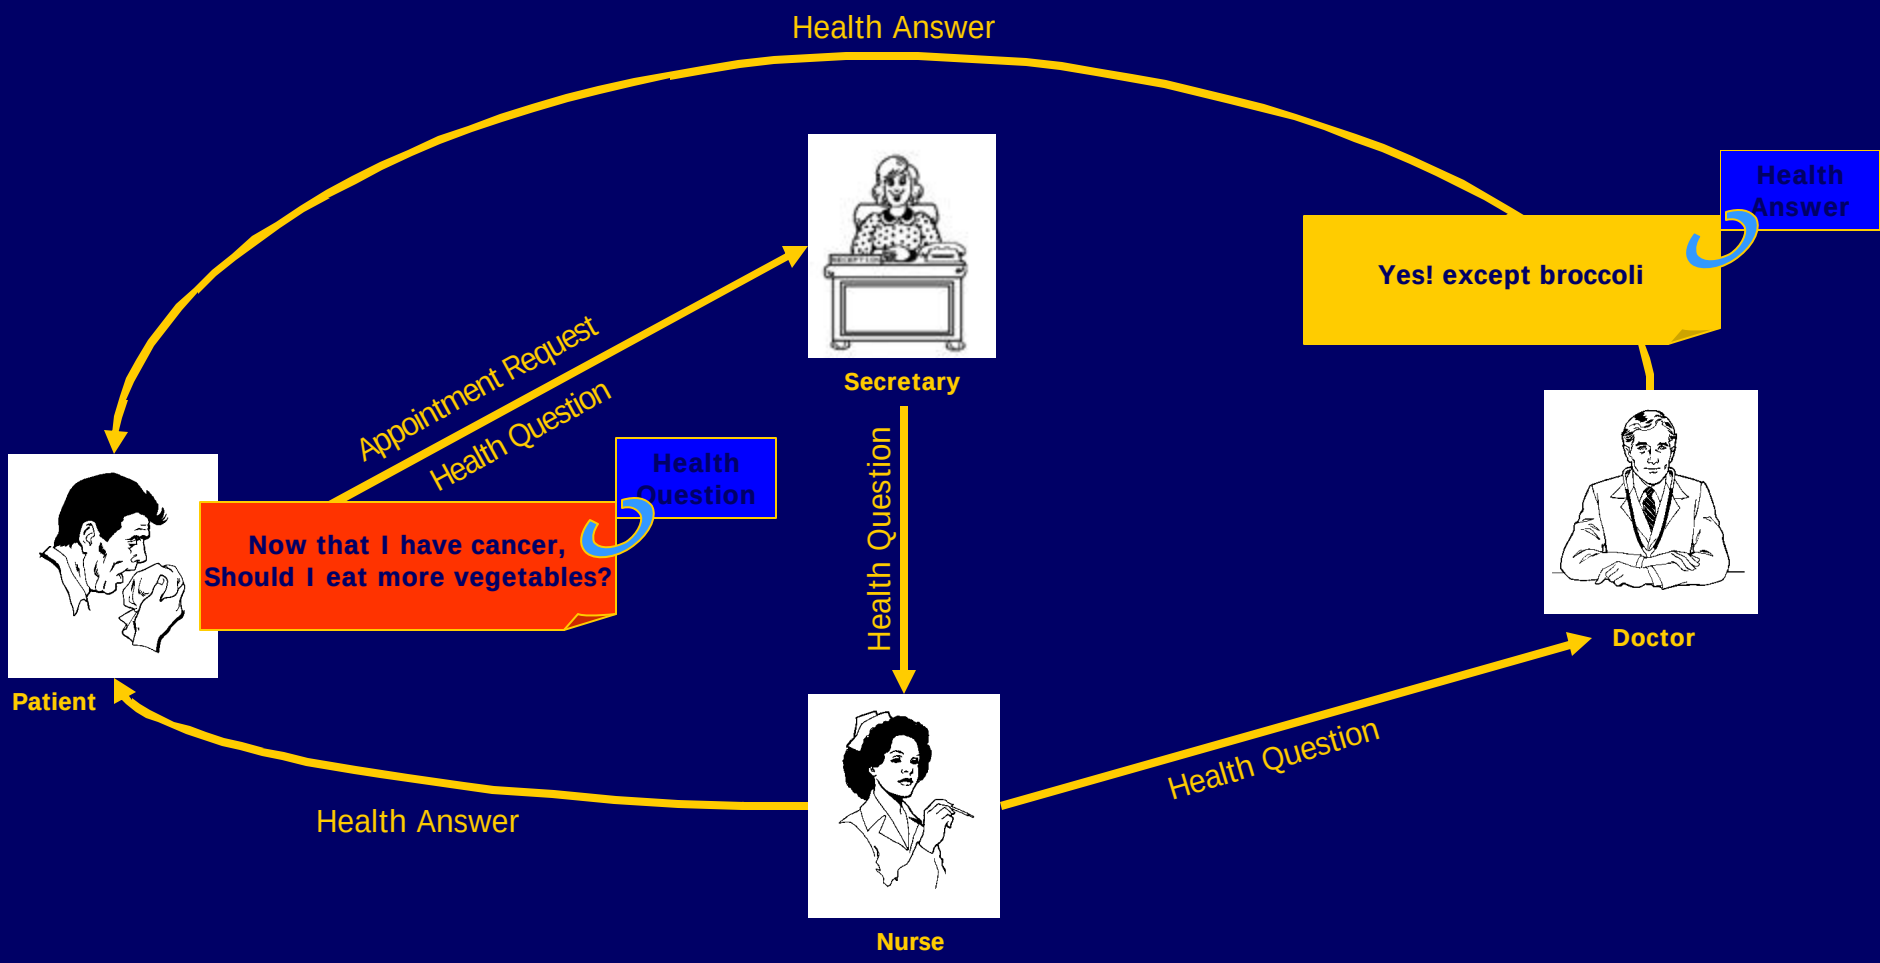

MyHealth@Vanderbilt Improved

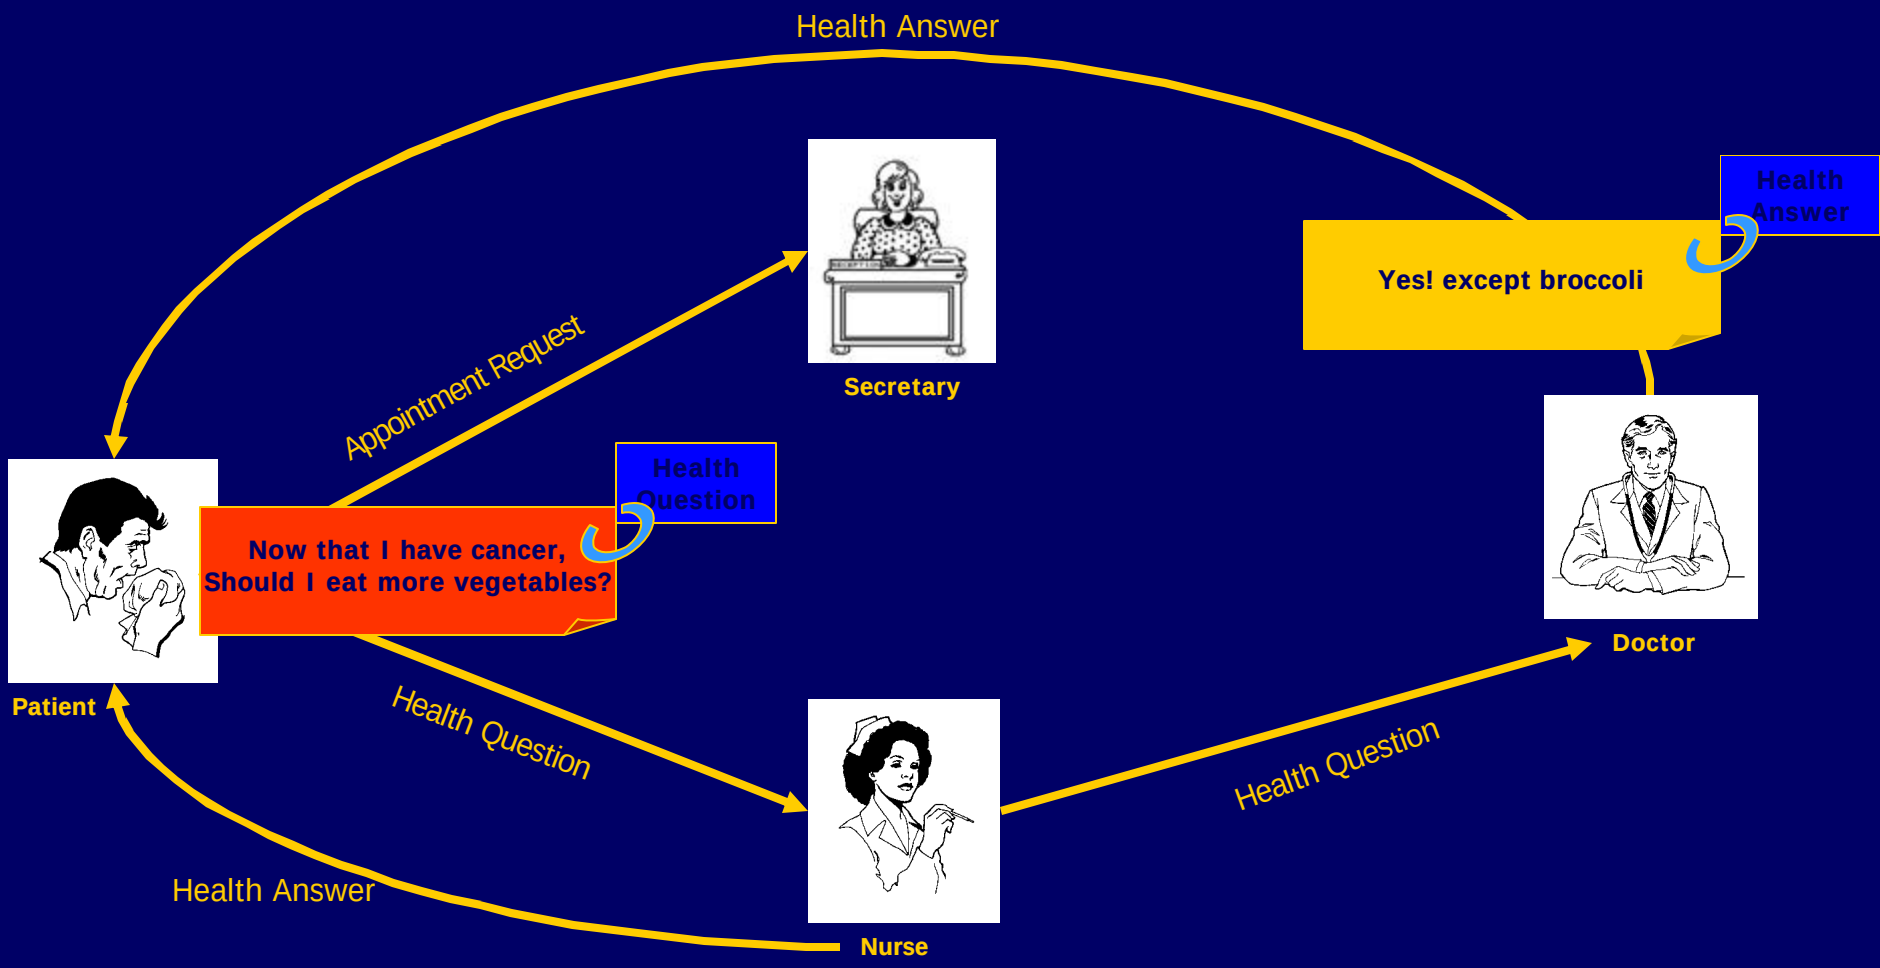

Outline

1. Motivating Example
2. Framework
   - Model
   - Logic of Privacy and Utility
3. Workflows and Responsibility
4. Algorithmic Results
   - Workflow Design assuming agents responsible
     - Abstract workflows
   - Auditing logs when agents irresponsible
     - Only graph-based workflows
5. Conclusions

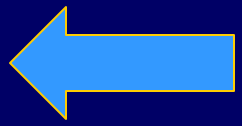

Auditing Results

- Definitions
  - Policy compliance, locally compliant
  - Causality, accountability
- Design of audit log
- Algorithms
  - Finding agents accountable for locally-compliant policy violation in graph-based workflows using audit log
  - Finding agents who act irresponsibly using audit log
- Algorithms use oracle:
  - ***O***(msg) = contents(msg)
  - Minimize number of oracle calls

Policy compliance/violation

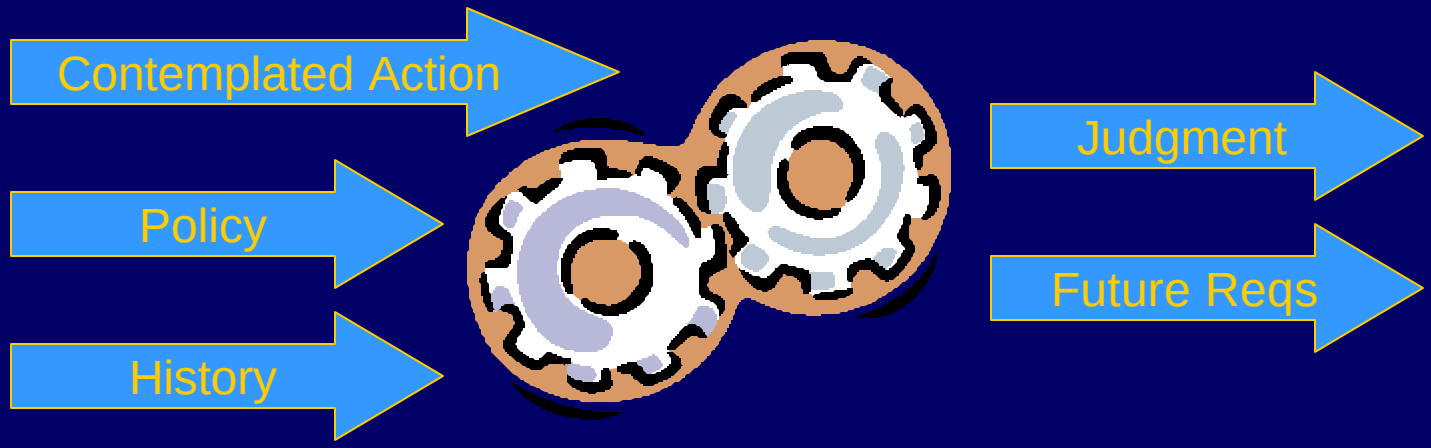

- Strong compliance [BDMN2006]
  - Action does not violate current policy requirements
  - Future policy requirements after action can be met
- Locally compliant policy
  - Agents can determine strong compliance based on their local view of history

Causality

- Lamport Causality [1978]

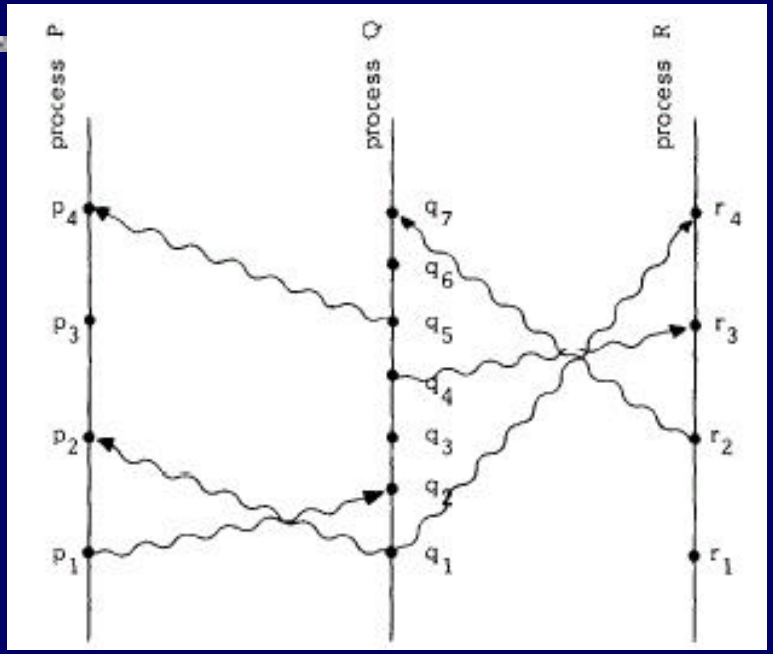

Definition. The relation "$\rightarrow$" on the set of events of a system is the smallest relation satisfying the following three conditions: (1) If a and b are events in the same process, and a comes before b, then $a \rightarrow b$. (2) If a is the sending of a message by one process and b is the receipt of the same message by another process, then $a \rightarrow b$. (3) If $a \rightarrow b$ and $b \rightarrow c$ then $a \rightarrow c$. Two distinct events a and b are said to be *concurrent* if $a \not\rightarrow b$ and $b \not\rightarrow a$.

Accountability & Audit Log

- Accountability
  - Causality + Irresponsibility
- Audit log design
  - Records all *Send(p,q,m)* and *Receive(p,q,m)* events executed
  - Maintains causality structure
    - O(1) operation per event logged

Auditing Algorithm

- Goal

  Find agents accountable for a policy violation

- Algorithm(Audit log A, Violation v)

  1. Construct G, the causality graph for v in A
  2. Run BFS on G.

     At each Send(p, q, m) node, check if tags(m) = O(m). If not, and p missed a tag, output p as accountable

- Theorem:

  - The algorithm outputs at least one accountable agent for every violation
    - of a locally compliant policy in an audit log
    - of a graph-based workflow that achieves the policy in the responsible model

Proof Idea

- Causality graph G includes all accountable agents
  - Accountability = Causality + Irresponsibility
- There is at least one irresponsible agent in G
  - Policy is satisfied if all agents responsible
  - Policy is locally compliant
- In graph-based workflows, safety responsibilities violated only by mistagging
  - *O*(msg) = tags(msg) check identifies all irresponsible actions

MyHealth Example

1. Policy violation:

   Secretary Candy receives health-question mistagged as appointment-request

2. Construct causality graph G and search backwards using BFS

   Candy received message m from Patient Jorge.

   - *O*(m) = health-question, but tags(m) = appointment-request.
   - Patient responsible for health-question tag.
   - Jorge identified as accountable

Conclusions

1. Framework
   - Concurrent game model
   - Logic of Privacy and Utility
     - Temporal logic (LTL, ATL*)
2. Business Process as Workflow
   - Role-based responsibility for human and mechanical agents
3. Algorithmic Results
   - Workflow design assuming agents responsible
     - Privacy, utility decidable (model-checking)
     - Minimal disclosure workflow constructible

     Autom-ated
   - Auditing logs when agents irresponsible
     - From policy violation to accountable agents
     - Finding irresponsible agents

     Using oracle

Thanks

Questions?

Local communication game

- Quotient structure under invisible actions, G_p
  - States:

    Smallest equivalence relation

    $K_1 \sim_p K_2$ if $K_1 \stackrel{a}{\Rightarrow} K_2$ and a is invisible to p
  - Actions:

    $[K] \stackrel{a}{\rightarrow} [K']$ if there exists K_1 in $[K]$ and K_2 in $[K']$ s.t. $K_1 \stackrel{a}{\Rightarrow} K_2$
- Lemma: For all LTL formulas φ visible to p, $G_p \models \langle\langle p \rangle\rangle \varphi$ implies $G \models \langle\langle p \rangle\rangle \varphi$

Refinement and Combination

- Policy refinement
  - Basic policy relation
  - Does hospital policy enforce HIPAA?
- P_1 refines P_2 if $P_1 \rightarrow P_2$
  - Requires careful handling of attribute inheritance
  - PSPACE decidable
- Combination becomes logical conjunction
  - Defined in terms of refinement

Related Languages

Model	Sender	Recipient	Subject	Attributes	Past	Future	Combination
RBAC	Role	Identity	×	×	×	×	•
XACML	Flexible	Flexible	Flexible	o	×	o	•
EPAL	Fixed	Role	Fixed	•	×	o	×
P3P	Fixed	Role	Fixed	•	o	×	o
LPU	Role	Role	Role	•	•	•	•

- Legend:
  - × unsupported
  - o partially supported
  - • fully supported
- LPU fully supports attributes, combination, temporal conditions

Why Not Use P3P?

- Different application
  - P3P understood by web browsers
  - LPU intended for internal policy enforcement
- Not expressive enough
  - P3P cannot express HIPAA, GLBA, COPPA
  - Each policy only has one sender and one subject
  - Missing temporal conditions; only has simple opt-in / opt-out

Structure of Attributes