From Qualitative to Quantitative Proofs of Security Properties Using First-Order Conditional Logic

Joseph Y. Halpern*
Cornell University
Dept. of Computer Science
Ithaca, NY 14853
halpern@cs.cornell.edu
http://www.cs.cornell.edu/home/halpern

Abstract

A first-order conditional logic is considered, with semantics given by a variant of ϵ-*semantics* (Adams 1975; Goldszmidt & Pearl 1992), where $\varphi {\rightarrow} \psi$ means that $\Pr(\psi \mid \varphi)$ approaches 1 *super-polynomially*—faster than any inverse polynomial. This type of convergence is needed for reasoning about security protocols. A complete axiomatization is provided for this semantics, and it is shown how a qualitative proof of the correctness of a security protocol can be automatically converted to a quantitative proof appropriate for reasoning about concrete security.

1 Introduction

Security protocols, such as key-exchange and key-management protocols, are short, but notoriously difficult to prove correct. Flaws have been found in numerous protocols, ranging from the the 802.11 Wired Equivalent Privacy (WEP) protocol used to protect link-layer communications from eavesdropping and other attacks (Borisov, Goldberg, & Wagner 2001) to standards and proposed standards for Secure Socket Layer (Wagner & Schneier 1996; Mitchell, Shmatikov, & Stern 1998) to Kerberos (Bella & Paulson 1998). Not surprisingly, a great deal of effort has been devoted to proving the correctness of such protocols. There are two largely disjoint approaches. The first essentially ignores the details of cryptography by assuming perfect cryptography (i.e., nothing encrypted can ever be decrypted without the encryption key) and an adversary that controls the network. By ignoring the cryptography, it is possible to give a more qualitative proof of correctness, using logics designed for reasoning about security protocols. Indeed, this approach has enabled axiomatic proofs of correctness and model checking of proofs (see, for example, (Mitchell, Mitchell, & Stern 1997; Paulson 1994)). The second approach applies the tools of modern cryptography to proving correctness, using more quantitative arguments. Typically it is shown that, given some security parameter k (where k may be, for example, the length of the key used) an adversary whose running time is polynomial in k has a negligible probability of breaking the security, where "negligible" means "less than any inverse polynomial function of k" (see, for example, (Bellare, Canetti, & Krawczyk 1998; Goldreich 2001)).

There has been recent work on bridging the gap between these two approaches, with the goal of constructing a logic that can allow reasoning about quantitative aspects of security protocols while still being amenable to mechanization. This line of research started with the work of Abadi and Rogaway 2000. More recently, Datta et al. 2005 showed that by giving a somewhat nonstandard semantics to their first-order *Protocol Composition Logic* (Datta *et al.* 2007), it was possible to reason about many features of the computational model. In this logic, an "implication" of the form $\varphi \supset B$ is interpreted as, roughly speaking, the probability of B given φ is high. For example, a statement like `secret encrypted` $\supset$ `adversary does not decrypt the secret` says "with high probability, if the secret is encrypted, the adversary does not decrypt it". While the need for such statements should be clear, the probabilistic interpretation used is somewhat unnatural, and no axiomatization is provided by Datta et al. 2005 for the $\supset$ operator (although some sound axioms are given that use it).

The interpretation of $\supset$ is quite reminiscent of one of the interpretations of $\rightarrow$ in conditional logic, where $\varphi \rightarrow \psi$ can be interpreted as "typically, if φ then ψ" (Kraus, Lehmann, & Magidor 1990). Indeed, one semantics given to $\rightarrow$, called ϵ-*semantics* (Adams 1975; Goldszmidt & Pearl 1992), is very close in spirit to that used in (Datta *et al.* 2005); this is particularly true for the formulation of ϵ-semantics given by Goldszmidt, Morris, and Pearl 1993. In this formulation, a formula $\varphi \rightarrow \psi$ is evaluated with respect to a sequence $(\Pr_1, \Pr_2, \ldots)$ of probability measures (*probability sequence*, for short): it is true if, roughly speaking, $\lim_{n\to\infty} \Pr_n(\psi \mid \varphi) = 1$ (where $\Pr_k(\psi \mid \varphi)$ is taken to be 1 if $\Pr_k(\varphi) = 1$). This formulation is not quite strong enough for some security-related purposes, where the standard is *super-polynomial* convergence, that is, convergence faster than any inverse polynomial. To capture such convergence, we can take $\varphi \rightarrow \psi$ to be true with respect to this probability sequence if, for all polynomials p, there exists n^* such that, for all $n \geq n^*$, $\Pr_n(\psi \mid \varphi) \geq 1 - 1/p(n)$. (Note that this implies that $\lim_{n\to\infty} \Pr_n(\psi \mid \varphi) = 1$.) In a companion paper, it is shown that reinterpreting $\rightarrow$ in this

*Supported in part by NSF under under grants ITR-0325453 and IIS-0534064, and by AFOSR under grant FA9550-05-1-0055.

way gives an elegant, powerful variant of the logic considered in (Datta *et al.* 2005), which can be used to reason about security protocols of interest.

While it is already a pleasant surprise that conditional logic provides such a clean approach to reasoning about security, using conditional logic has two further significant advantages, which are the subject of this paper. The first is that, as I show here, the well-known complete axiomatization of conditional logic with respect to ϵ-semantics continues to be sound and complete with respect to the super-polynomial semantics for $\rightarrow$; thus, the axioms form a basis for automated proofs. The second is that the use of conditional logic allows for a clean transition from qualitative to quantitative arguments. To explain these points, I need to briefly recall some well-known results from the literature.

As is well known, the *KLM properties* (Kraus, Lehmann, & Magidor 1990) (see Section 2) provide a sound and complete axiomatization for reasoning about $\rightarrow$ formulas with respect to ϵ-semantics (Geffner 1992). More precisely, if Δ is a collection of formulas of the form $\varphi' \rightarrow \psi'$, then Δ *(ϵ-)entails* $\varphi \rightarrow \psi$ (that is, for every probability sequence $\mathcal{P}$, if every formula in Δ is true in $\mathcal{P}$ according to ϵ semantics, then so is $\varphi \rightarrow \psi$), then $\varphi \rightarrow \psi$ is provable from Δ using the KLM properties. This result applies only when Δ is a collection of $\rightarrow$ formulas. Δ cannot include negations or disjunctions of $\rightarrow$ formulas. *Conditional logic* extends the KLM framework by allowing Boolean combinations of $\rightarrow$ statements. A sound and complete axiomatization of propositional conditional logic with semantics given by what are called preferential structures was given by Burgess 1981; Friedman and Halpern 2001 proved it was also sound and complete for ϵ-semantics.

Propositional conditional logic does not suffice for reasoning about security. The logic of (Datta *et al.* 2005) is first-order; quantification is needed to capture important properties of security protocols. A sound and complete axiomatization for the language of first-order conditional logic, denoted $\mathcal{L}_C$, with respect to ϵ-semantics is given by Friedman, Halpern, and Koller 2000. The first major result of this paper shows a conditional logic formula φ is satisfiable in some model M with respect to ϵ-semantics iff it is satisfiable in some model M' with respect to the super-polynomial semantics. It follows that all the completeness results for ϵ-semantics apply without change to the super-polynomial semantics.

I then consider the language $\mathcal{L}_C^0$ which essentially consists of universal $\rightarrow$ formulas, that is, formulas of the form $\forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$, where φ and ψ are first-order formulas. As in the KLM framework, there are no nested $\rightarrow$ formulas or negated $\rightarrow$ formulas. The second major result of this paper is to provide a sound and complete axiomatization that extends the KLM properties for reasoning abut when a collection of formulas in $\mathcal{L}_C^0$ entails a formula in $\mathcal{L}_C^0$.

It might seem strange to be interested in an axiomatization for $\mathcal{L}_C^0$ when there is already a sound and complete axiomatization for the full language $\mathcal{L}_C$. However, $\mathcal{L}_C^0$ has some significant advantages. In reasoning about concrete security, asymptotic complexity results do not suffice; more detailed information about security guarantees is needed. For example, we may want to prove that an SSL server that supports 1,000,000 sessions using 1024 bit keys has a probability of 0.999999 of providing the desired service without being compromised. I show how to convert a qualitative proof of security in the language $\mathcal{L}_C^0$, which provides only asymptotic guarantees, to a quantitative proof. Moreover, the conversion shows exactly how strong the assumptions have to be in order to get the desired 0.999999 level of security. Such a conversion is not possible with $\mathcal{L}_C$.

This conversion justifies reasoning at the qualitative level. A qualitative proof can be constructed without worrying about the details of the numbers, and then automatically converted to a quantitative proof for the desired level of security.

In the next section, I review the syntax and semantics of conditional logic, with an emphasis on ϵ semantics, and show how it can be modified to deal with the super-polynomial convergence that is more appropriate for reasoning about security. In Section 3, I provide axioms and inference rules for both qualitative and quantitative reasoning.

2 First-Order Conditional Logic

I review the syntax and semantics of first-order conditional logic here. Although I focus on first-order conditional logic here, it is straightforward to specialize all the definitions and results to the propositional case, so I do not discuss the propositional case further.

The syntax of first-order conditional logic is straightforward. Fix a finite first-order *vocabulary* $\mathcal{T}$ consisting, as usual, of function symbols, predicate symbols, and constants. Starting with atomic formulas of first-order logic over the vocabulary $\mathcal{T}$, more complicated formulas are formed by closing off under the standard truth-functional connectives (i.e., $\wedge$,$\vee$, $\neg$, and $\Rightarrow$), first-order quantification, and the binary modal operator $\rightarrow$. Thus, a typical formula is $\forall x(P(x) \rightarrow \exists y(Q(x,y) \rightarrow R(y)))$. Let $\mathcal{L}_C(\mathcal{T})$ be the resulting language. Let $\mathcal{L}^{fo}(\mathcal{T})$ be the pure first-order fragment of $\mathcal{L}_C(\mathcal{T})$, consisting of $\rightarrow$-free formulas. Let $\mathcal{L}_C^0(\mathcal{T})$ consist of all formulas in $\mathcal{L}_C(\mathcal{T})$ of the form $\forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$, where φ and ψ are in $\mathcal{L}^{fo}$. (I henceforth omit the $\mathcal{T}$ unless it is necessary for clarity.) Note that $\mathcal{L}_C^0$ does not include negations of $\rightarrow$ formulas or conjunctions of $\rightarrow$ formulas. While not having conjunctions does not really impair the expressive power of $\mathcal{L}_C^0$ (since we will be interested in sets of $\mathcal{L}_C^0$ formulas, where a set can be identified with the conjunction of the formulas in the set), the lack of negation does.

I give two semantics to formulas in $\mathcal{L}_C(\mathcal{T})$. In both semantics, the truth of formulas is defined with respect to *PS structures*. A PS structure is a tuple $M = (D, W, \pi, \mathcal{P})$, where D is a domain, W is a set of worlds, π is an *interpretation*, which associates with each predicate symbol (resp., function symbol, constant symbol) in $\mathcal{T}$ and world $w \in W$ a predicate (resp., function, domain element) of the right arity, and $\mathcal{P} = (\Pr_1, \Pr_2, \ldots)$ is a probability sequence. As usual, a *valuation* V associates with each variable x an element $V(x) \in D$.

Given a valuation V and structure M, the semantics of $\wedge$, $\neg$, $\Rightarrow$, and $\forall$ is completely standard. In particular, the truth of a first-order formula in $\mathcal{L}^{fo}$ in a world w, written $(M, V, w) \models \varphi$, is determined as usual. For $\varphi \in \mathcal{L}^{fo}$, let

$[\![\varphi]\!]_{M,V} = \{w : (M, V, w) \models \varphi\}$. If φ is a closed formula, so that its truth does not depend on the valuation, I occasionally write $[\![\varphi]\!]_M$ rather than $[\![\varphi]\!]_{M,V}$. I write $(M, V) \models \varphi$ if $(M, V, w) \models \varphi$ for all worlds w. The truth of an $\rightarrow$ formula does not depend on the world, but only on the structure M.

$$(M, V) \models \varphi{\rightarrow}\psi \text{ if } \lim_{n\rightarrow\infty} \Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) = 1,$$

where $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V})$ is taken to be 1 if $\Pr_n([\![\varphi]\!]_{M,V}) = 0$.

I also consider an alternative semantics that gives super-polynomial convergence.

$(M, V) \models^{sp} \varphi{\rightarrow}\psi$ if for all k,
there exists some $n^* \geq 0$ such that, for all $n \geq n^*$,
$\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - (1/n^k)$.$\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - (1/p(n))$.

Note that if $(M, V) \models^{sp} \varphi{\rightarrow}\psi$ and $p(n)$ is a polynomial whose leading coefficient is positive, then $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - (1/p(n))$ for sufficiently large n.

As usual, I write $M \models \varphi$ if $(M, V) \models \varphi$ for all valuations V, and $\mathcal{M} \models \varphi$ if $M \models \varphi$ for all PS structures in a set $\mathcal{M}$, and similarly with $\models$ replaced by $\models^{sp}$.

3 Axioms for qualitative and quantitative reasoning

In this section, I start by showing that qualitative reasoning for both $\models$ and $\models^{sp}$ is characterized by the same axiom system. I then provide a complete axiomatization for $\mathcal{L}_C^0$. Finally, I consider quantitative conditional logic. In the axioms, it is convenient to use $N\varphi$ as an abbreviation for $\neg\varphi{\rightarrow}\mathit{false}$. Note that if φ is a closed formula, then $M \models N\varphi$ iff, for some n^*, $\Pr_n([\![\varphi]\!]_M) = 0$ for all $n \geq n^*$, and similarly with $\models$ replaced by $\models^{sp}$. Thus, $N\varphi$ can be read as saying "φ is almost surely true".

3.1 Qualitative Reasoning

As was mentioned in the introduction, Friedman, Halpern, and Koller 2000 provide a complete axiomatization AX_C for $\mathcal{L}_C$ with respect to $\models$. For security applications, a generalization of their result is needed, where it is possible to restrict to models where all worlds satisfy a particular first-order theory Λ. (We can think of Λ as describing first-order properties of the security protocol being analyzed.)

Let $\vdash_\Lambda$ denote provability in first-order logic given the axioms in the theory Λ. Let AX_C^Λ consist of the following axioms and rules:

Λ-AX. φ, if $\varphi \in \mathcal{L}^{fo}$ and $\vdash_\Lambda \varphi$.

C0. All substitution instances of propositional tautologies.

C1. $\varphi{\rightarrow}\varphi$.

C2. $((\varphi{\rightarrow}\psi_1) \wedge (\varphi{\rightarrow}\psi_2)) \Rightarrow (\varphi{\rightarrow}(\psi_1 \wedge \psi_2))$.

C3. $((\varphi_1{\rightarrow}\psi) \wedge (\varphi_2{\rightarrow}\psi)) \Rightarrow ((\varphi_1 \vee \varphi_2){\rightarrow}\psi)$.

C4. $((\varphi_1{\rightarrow}\varphi_2) \wedge (\varphi_1{\rightarrow}\psi)) \Rightarrow ((\varphi_1 \wedge \varphi_2){\rightarrow}\psi)$.

C5. $[(\varphi{\rightarrow}\psi) \Rightarrow N(\varphi{\rightarrow}\psi)] \wedge [\neg(\varphi{\rightarrow}\psi) \Rightarrow N\neg(\varphi{\rightarrow}\psi)]$.

C6. $\neg(\mathit{true}{\rightarrow}\mathit{false})$.

F1. $\forall x\varphi \Rightarrow \varphi[x/t]$, where t is *substitutable* for x in the sense discussed below and $\varphi[x/t]$ is the result of substituting t for all free occurrences of x in φ (see (Enderton 1972) for a formal definition).

F2. $\forall x(\varphi \Rightarrow \psi) \Rightarrow (\forall x\varphi \Rightarrow \forall x\psi)$.

F3. $\varphi \Rightarrow \forall x\varphi$ if x does not occur free in φ.

F4. $x = y \Rightarrow (\varphi_1 \Rightarrow \varphi_2)$, where φ_1 is quantifier-free and φ_2 is obtained from φ_1 by replacing zero or more occurrences of x in φ_1 by y.

F5. $x \neq y \Rightarrow N(x \neq y)$.

MP. From φ and $\varphi \Rightarrow \psi$ infer ψ.

Gen. From φ infer $\forall x\varphi$.

R1. From $\varphi_1 \Leftrightarrow \varphi_2$ infer $\varphi_1{\rightarrow}\psi \Leftrightarrow \varphi_2{\rightarrow}\psi$.

R2. From $\psi_1 \Rightarrow \psi_2$ infer $\varphi{\rightarrow}\psi_1 \Rightarrow \varphi{\rightarrow}\psi_2$.

The axiom system AX_C of (Friedman, Halpern, & Koller 2000) does not have Λ-AX (this is needed to incorporate the theory Λ) and includes an axiom $x = x$ that follows from Λ-AX; otherwise, the axiom systems are identical. As observed in (Friedman, Halpern, & Koller 2000), the "positive" version of F5, $x = y \Rightarrow N(x = y)$, is also sound. It is not included in the axiomatization because it is provable from the other axioms.

It remains to explain the notion of "substitutable" in F1. Clearly a term t with free variables that might be captured by some quantifiers in φ cannot be substituted for x; for example, while $\forall x\exists y(x \neq y)$ is true as long as the domain has at least two elements, the result of substituting y for x is $\exists y(y \neq y)$, which is surely false. In the case of first-order logic, it suffices to define "substitutable" so as to make sure this does not happen (see (Enderton 1972) for details). However, in modal logics such as this one, more care must be taken. In general, terms cannot be substituted for universally quantified variables in a modal context, since terms are not in general *rigid*; that is, they can have different interpretations in different worlds. To understand the impact of this, consider the formula $\forall x(\neg NP(x)) \Rightarrow \neg NP(c)$ (where P is a unary predicate and c is a constant). This formula is not valid in PS structures. For example, consider a PS structure with two worlds w_1 and w_2, and a domain with two elements d_1 and d_2. Suppose that in world w_1, $P(d_1)$ holds, $P(d_2)$ does not, and c is interpreted as d_1, while in world w_2, $P(d_2)$ holds, $P(d_1)$ does not, and c is interpreted as d_2. Then it is easy to see that $NP(c)$ holds in both worlds, but $NP(x)$ holds in only one world, no matter how x is interpreted. If $Pr_n(w_1) = Pr_n(w_2) = 1/2$ for all n, then $M \models NP(c)$, while $M \models \forall x(\neg NP(x))$. Thus, if φ is a formula that has occurrences of $\rightarrow$, then the only terms that are considered substitutable for x in φ are other variables.

It is interesting to contrast R1, R2, and C4. While R2 allows a formula ψ_1 on the right-hand side of $\rightarrow$ to be replaced by a weaker formula ψ_2 (that is, a formula such that that $\psi_1 \Rightarrow \psi_2$ is provable), R1 just allows a formula φ_1 on the left-hand side of $\rightarrow$ to be replaced by an equivalent formula φ_2, rather than a stronger formula. C4 allows the replacement of a formula φ_1 on the left-hand side by a stronger formula, $\varphi_1 \wedge \varphi_2$, but only if $\varphi_1{\rightarrow}\varphi_2$ holds. Intuitively, this says that if φ_2 and ψ each almost always hold given φ_1, then ψ almost always holds given both φ_1 and φ_2. Monotonicity does not hold in general. That is, if $\varphi_1 \Rightarrow \varphi_2$ is provable and $\varphi_2 \rightarrow \psi$ holds, then $\varphi_1 \rightarrow \psi$ does not necessarily hold. For a simple counterexample, it is not the case that if $true \rightarrow \psi$ holds then $\neg\psi \rightarrow \psi$ holds. If $\psi(x)$ states that x cannot break the encryption, it seems reasonable to expect that, almost always, x cannot break the encryption ($true \rightarrow \psi(x)$), but it surely is not the case that x cannot break the encryption given that x can break it.

The remaining axioms and rules are easy to explain. In particular, C2 says that if both ψ_1 and ψ_2 almost always hold given φ, then so does $\psi_1 \wedge \psi_2$, while C3 allows reasoning by cases: if ψ almost always holds given each of φ_1 and φ_2, then it almost always holds given their disjunction.

I want to show that AX_C^Λ is also sound and complete for the $\models^{sp}$ semantics. The key step in doing that is to show that a formula is satisfiable with respect to the $\models$ semantics iff it is satisfiable with respect to the $\models^{sp}$ semantics.

Theorem 3.1: *If $M = (D, W, \pi, \mathcal{P})$ is a PS structure and D is countable, then there exists a probability sequence $\mathcal{P}'$ such that, for all valuations V, $(M, V) \models \varphi$ iff $(M', V) \models^{sp} \varphi$, where $M' = (D, W, \pi, \mathcal{P}')$.*

Proof: Suppose that $M = (D, W, \pi, \mathcal{P})$, where $D = \{d_1, d_2, \ldots\}$ (D may be finite), and $\mathcal{P} = (\Pr_1, \Pr_2, \ldots)$. Let $L = (\varphi_1 \rightarrow \psi_1, \varphi_2 \rightarrow \psi_2, \ldots)$ be a list of all formulas of the form $\varphi' \rightarrow \psi'$ in $\mathcal{L}_C$ with the property that if $(M, V') \models \neg(\varphi' \rightarrow \psi')$ for some valuation V', then $\varphi' \rightarrow \psi'$ appears infinitely often in L. Suppose that the set of variables is $\{x_1, x_2, \ldots\}$. (I am implicitly assuming that the set of variables is countable, as is standard.) Let $\mathcal{V}_n$ be the set of valuations V such that $V(x_i) \in \{d_1, \ldots, d_n\}$ for $i = 1, \ldots, n$ and $V(x_m) = d_1$ for all $m > n$. Given a valuation V' and a formula $\varphi \in \mathcal{L}_C$, there exists n such that, for all free variables x in φ, $x \in \{x_1, \ldots, x_n\}$ and $V'(x) \in \{d_1, \ldots, d_n\}$. Thus, $(M, V') \models \varphi$ for some valuation V' iff $(M, V') \models \varphi$ for some valuation $V' \in \mathcal{V}_n$. Suppose that the elements of $\mathcal{V}_n$ are $V_1^n, \ldots, V_{|\mathcal{V}_n|}^n$.

Since $\mathcal{V}_n$ is finite, there is a subsequence $\mathcal{P}' = (\Pr'_{11}, \ldots \Pr_{1|\mathcal{V}_1|}, \Pr'_{21}, \ldots, \Pr_{2|\mathcal{V}_2|}, \ldots)$ of $\mathcal{P}$ with the following properties, for $1 \leq m \leq |\mathcal{V}_n|$:

$$\begin{array}{l}\text{for all } j \leq n \text{ and } V' \in \mathcal{V}_n, \text{ if } (M, V') \models \varphi_j \rightarrow \psi_j, \\ \text{then } \Pr'_{nm}([\![\psi_j]\!]_{M,V'} \mid [\![\varphi_j]\!]_{M,V'}) \geq 1 - 1/n^n;\end{array} \tag{1}$$

$$\begin{array}{l}\text{if } (M, V_m^n) \models \neg(\varphi_n \rightarrow \psi_n), \text{ then} \\ \Pr'_{nm}([\![\psi_n]\!]_{M,V_m^n} \mid [\![\varphi_n]\!]_{M,V_m^n}) < 1 - 1/k, \text{ where } k \\ \text{is the smallest integer such that, for infinitely many} \\ \text{indices } h, \Pr_h([\![\psi_n]\!]_{M,V_m^n} \mid [\![\varphi_n]\!]_{M,V_m^n}) < 1 - 1/k.\end{array} \tag{2}$$

(There must be such a k, since $\lim_{h \rightarrow \infty} \Pr_h([\![\psi_n]\!]_{M,V_m^n} \mid [\![\varphi_n]\!]_{M,V_m^n}) \neq 1$.)

Let $M' - (D, W, \pi, \mathcal{P}')$. I now prove that $(M, V) \models \varphi$ iff $(M', V) \models^{sp} \varphi$ for all valuations V and formulas $\varphi \in \mathcal{L}_C$ by a straightforward induction on the structure of φ. If φ is an atomic formula, this is immediate, since M and M' differ only in their probability sequences. All cases but the one where φ has the form $\varphi' \rightarrow \psi'$ follow immediately from the induction hypothesis. If φ has the form $\varphi' \rightarrow \psi'$, first suppose that $(M, V) \models \varphi' \rightarrow \psi'$. Fix a polynomial p. There must exist some n^* such that (a) for all free variables x in φ' or ψ', $x \in \{x_1, \ldots, x_{n^*}\}$ and $V(x) \in \{d_1, \ldots d_{n^*}\}$, (b) $p(n) < 1/n^n$ for all $n \geq n^*$, and (c) $\varphi' \rightarrow \psi'$ is among the first n^* formulas in L. It follows from (a) that for all $n \geq n^*$, there exists some $V' \in \mathcal{V}_N$ such that V' and V agree on all the free variables in $\varphi' \rightarrow \psi'$. It then follows from (b), (c), and (1) that, for all $n \geq n^*$ and $1 \leq m \leq |\mathcal{V}_n|$, $\Pr'_{nm}([\![\psi']\!]_{M,V} \mid [\![\varphi']\!]_{M,V}) \geq 1 - 1/p(n)$. Thus, $(M, V) \models^{sp} \varphi' \rightarrow \psi'$.

If $(M, V) \models \neg(\varphi' \rightarrow \psi')$, there must be some minimal k such that $\Pr_h([\![\psi']\!]_{M,V} \mid [\![\varphi']\!]_{M,V}) < 1 - 1/k$ for infinitely many indices h. Since $\varphi' \rightarrow \psi'$ occurs infinitely often in L, it easily follows from (2) that, for infinitely many values of n and h, $\Pr'_{nh}([\![\psi']\!]_{M,V} \mid [\![\varphi']\!]_{M,V}) < 1 - 1/k$. Let $p(n) = k$ (so $p(n)$ is a constant function). It follows that $\Pr'_{nh}([\![\psi']\!]_{M,V} \mid [\![\varphi']\!]_{M,V}) < 1 - 1/p(n)$ for infinitely many values of n and h. Thus, $(M, V) \models^{sp} \neg(\varphi' \rightarrow \psi')$. This completes the proof. ∎

Let $\mathcal{PS}(\Lambda)$ consist of all PS structures M where every world satisfies Λ.

Theorem 3.2: *AX_C^Λ is a sound and complete axiomatization for $\mathcal{PS}(\Lambda)$ with respect to both $\models$ and $\models^{sp}$. That is, the following are equivalent for all formulas in $\mathcal{L}_C(\mathcal{T})$:*

(a) $AX_C^\Lambda \vdash \varphi$;

(b) $\mathcal{PS}(\Lambda) \models \varphi$;

(c) $\mathcal{PS}(\Lambda) \models^{sp} \varphi$.

Proof: The equivalence of parts (a) and (b) for the case that $\Lambda = \emptyset$ is proved in Theorem 5.2 of (Friedman, Halpern, & Koller 2000). The same proof shows that the result holds for arbitrary Λ. To show that (a) implies (c), I must show that all the axioms are sound. The soundness of all the axioms and rules other than C2, C3, C4, and C5 is trivial. I consider each of these axioms in turn.

For C2, suppose that $M = (D, W, \pi, (\Pr_1, \Pr_2, \ldots))$ is a PS structure such that $M \models^{sp} \varphi \rightarrow \psi_1$ and $M \models^{sp} \varphi \rightarrow \psi_2$. Since $M \models^{sp} \varphi \rightarrow \psi_i$, $i = 1, 2$, given a positive polynomial p, there exists $n_1^*, n_2^* \geq 0$ such that, for all $n \geq n_i^*$, $\Pr_n([\![\psi_i]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - 1/2p(n)$, for $i = 1, 2$. For all $n \geq \max(n_1^*, n_2^*)$, $\Pr_n([\![\psi_i]\!] \mid [\![\varphi]\!]) \leq 1/2p(n)$. Thus, for $n \geq \max(n_1^*, n_2^*)$,

$$\begin{array}{rl} & \Pr_n([\![\psi_1 \wedge \psi_2]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \\ \geq & 1 - (\Pr_n([\![\psi_1]\!] \mid [\![\varphi]\!]) + \Pr_n([\![\psi_2]\!] \mid [\![\varphi]\!])) \\ \geq & 1 - \frac{1}{2p(n)} - \frac{1}{2p(n)} \\ = & 1 - \frac{1}{p(n)}. \end{array}$$

For C3, note that

$$\begin{array}{rl} & \Pr(A \mid B_1 \cup B_2) \\ = & \Pr((A \cap B_1 \cup A \cap B_2) \mid B_1 \cup B_2) \\ = & \Pr(A \cap B_1 \mid B_1 \cup B_2) + \Pr(A \cap B_2 \mid B_1 \cup B_2) \\ & - \Pr(A \cap B_1 \cap B_2 \mid B_1 \cup B_2) \\ = & \Pr(A \mid B_1) \times \Pr(B_1 \mid B_1 \cup B_2) \\ & + \Pr(A \mid B_2) \times \Pr(B_2 \mid B_1 \cup B_2) \\ & - \Pr(A \cap B_1 \cap B_2 \mid B_1 \cup B_2). \end{array} \tag{3}$$

Now suppose that $M \models^{sp} \varphi_1 \rightarrow \psi$ and $M \models^{sp} \varphi_2 \rightarrow \psi$. Given a positive polynomial p, as in the case of C2, there exist n_1^* and n_2^* such that, for all $n \geq n_i^*$, $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi_1]\!]_{M,V}) \geq 1 - 1/2p(n)$, for $i = 1, 2$. It easily follows from (3) that if $n \geq \max(n_1^*, n_2^*)$, then

$$\begin{array}{ll} & \Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ \geq & (1 - \frac{1}{2p(n)})\Pr_n([\![\varphi_1]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ & +(1 - \frac{1}{2p(n)})\Pr_n([\![\varphi_2]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ & -\Pr_n([\![\psi \wedge \varphi_1 \wedge \varphi_2]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ \geq & (1 - \frac{1}{2p(n)})\Pr_n([\![\varphi_1]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ & +(1 - \frac{1}{2p(n)})\Pr_n([\![\varphi_2]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ & -\Pr_n([\![\varphi_1 \wedge \varphi_2]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ \geq & (1 - \frac{1}{2p(n)})[\Pr_n([\![\varphi_1]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ & +\Pr_n([\![\varphi_2]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ & -\Pr_n([\![\varphi_1 \wedge \varphi_2]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V})] \\ & -\frac{1}{2p(n)}\Pr_n([\![\varphi_1 \wedge \varphi_2]\!]_{M,V} \mid [\![\varphi_1 \vee \varphi_2]\!]_{M,V}) \\ \geq & (1 - \frac{1}{2p(n)}) - \frac{1}{2p(n)} \\ = & 1 - \frac{1}{p(n)}. \end{array}$$

For C4, note that

$$\begin{array}{ll} & \Pr(A_1 \mid A_2 \cap B) \\ = & \Pr(A_1 \cap A_2 \mid B)/\Pr(A_2 \mid B) \geq \Pr(A_2 \cap A_2 \mid B), \end{array}$$

so the argument follows essentially the same lines as that for C2.

Finally, C5 follows easily from the fact that the truth of a formula of the form $\varphi \rightarrow \psi$ or $\neg(\varphi \rightarrow \psi)$ is independent of the world, and depends only on the probability sequence.

Finally, I must show that (c) implies (b). Suppose not. Then there exists a formula φ such that $\mathcal{PS}(\Lambda) \models^{sp} \varphi$ but $\mathcal{PS}(\Lambda) \not\models \varphi$. Thus, there exists $M \in \mathcal{PS}(\Lambda)$ and valuation V such that $(M, V) \not\models \varphi$. The proof in (Friedman, Halpern, & Koller 2000) shows that if a formula is satisfiable with respect to $\models$ at all, then it is satisfiable in a structure in $\mathcal{PS}(\Lambda)$ with a countable domain. Thus, without loss of generality, M has a countable domain. But then it immediately follows from Theorem 3.1 that $\mathcal{PS}(\Lambda) \not\models^{sp} \varphi$. ∎

I next completely characterize reasoning in $\mathcal{L}_C^0$. (Recall that $\mathcal{L}_C^0$ consists of all formulas of the form $\forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$, where φ and ψ are first-order formulas.) I start by considering the fragment $\mathcal{L}_C^-$ of $\mathcal{L}_C^0$ consisting of all formulas of the form $\varphi \rightarrow \psi$ where φ and ψ are closed first-order formulas. Thus, $\mathcal{L}_C^-$ does not allow $\rightarrow$ formulas to be universally quantified. Consider the following rules:

LLE. If $\vdash_\Lambda \varphi_1 \Leftrightarrow \varphi_2$, then from $\varphi_1 \rightarrow \psi$ infer $\varphi_2 \rightarrow \psi$ (left logical equivalence).

RW. If $\vdash_\Lambda \psi_1 \Rightarrow \psi_2$, then from $\varphi \rightarrow \psi_1$ infer $\varphi \rightarrow \psi_2$ (right weakening).

REF. $\varphi \rightarrow \varphi$ (reflexivity).

AND. From $\varphi \rightarrow \psi_1$ and $\varphi \rightarrow \psi_2$ infer $\varphi \rightarrow \psi_1 \wedge \psi_2$.

OR. From $\varphi_1 \rightarrow \psi$ and $\varphi_2 \rightarrow \psi$ infer $\varphi_1 \vee \varphi_2 \rightarrow \psi$.

CM. From $\varphi_1 \rightarrow \varphi_2$ and $\varphi_1 \rightarrow \psi$ infer $\varphi \wedge \varphi_2 \rightarrow \psi$ (cautious monotonicity).

This collection of rules has been called system $\mathbf{P}_\Lambda$ (Kraus, Lehmann, & Magidor 1990) or *the KLM properties*.[1] The rules are obvious analogues of axioms in AX_C^Λ. In particular, LLE is the analogue of R1, RW is the analogue of R2, REF is the analogue C1, AND is the analogue of C2, OR is the analogue of C3, and CM is the analogue of C4.

Given a collection Δ of $\rightarrow$ formulas, I write $\mathbf{P}_\Lambda \vdash \Delta \hookrightarrow \varphi \rightarrow \psi$ if $\varphi \rightarrow \psi$ can be derived from Δ using these rules. A *derivation from* Δ consists of a sequence of steps of the form $\Delta \hookrightarrow \varphi \rightarrow \psi$, where either (a) $\varphi \rightarrow \psi \in \Delta$, (b) $\varphi = \psi$ (which can be viewed as an application of the axiom REF), or (c) $\varphi \rightarrow \psi$ follows from previous steps by application of one of the rules in $\mathbf{P}_\Lambda$. All the rules above have the form "from $\varphi_1 \rightarrow \psi_1, \ldots \varphi_n \rightarrow \psi_i$' infer $\varphi \rightarrow \psi$"; this can be viewed as an abbreviation for the rule scheme "from $\Delta \hookrightarrow \varphi_1 \rightarrow \psi_i, \ldots \Delta \hookrightarrow \varphi_n \rightarrow \psi_n$ infer $\Delta \hookrightarrow \varphi \rightarrow \psi$", with the same Δ everywhere. Although, for all these rules, the set Δ is the same everywhere, later there will be rules where different sets Δ are involved. I write $(M, V) \models \Delta \hookrightarrow \varphi \rightarrow \psi$ if $(M, V) \models \varphi' \rightarrow \psi'$ for every formula $\varphi' \rightarrow \psi' \in \Delta$ implies that $(M, V) \models \varphi \rightarrow \psi$. (For a formula $\varphi \rightarrow \psi \in \mathcal{L}_C^-$, φ and ψ are closed, so $(M, V) \models \varphi \rightarrow \psi$ iff $M \models \varphi \rightarrow \psi$. However, in $\mathcal{L}_C^0$ there are open formulas, so the valuation V plays a role.) I write $\mathcal{PS}(\Lambda) \models \Delta \hookrightarrow \varphi \rightarrow \psi$ if $(M, V) \models \Delta \hookrightarrow \varphi \rightarrow \psi$ for all PS structures M and valuations V. As usual, a rule is said to be *sound* if it preserves truth (in this case, with respect to all (M, V)); that is, if all the antecedents hold with respect to (M, V), then so does the conclusion.

The following result is well known.

Theorem 3.3: *(Kraus, Lehmann, & Magidor 1990; Geffner 1992) If $\Delta \cup \{\varphi \rightarrow \psi\} \subseteq \mathcal{L}_C^-$, then $\mathbf{P}_\Lambda \vdash \Delta \hookrightarrow \varphi \rightarrow \psi$ iff $\mathcal{PS}(\Lambda) \models \Delta \hookrightarrow \varphi \rightarrow \psi$.*

I want to extend this result from $\mathcal{L}_C^-$ to $\mathcal{L}_C^0$, and to the $\models^{sp}$ semantics as well as the $\models$ semantics, so as to make it applicable to reasoning about security protocols. I actually extend it to $\mathcal{L}_C^0 \cup \mathcal{L}^{fo}$. A collection Δ of formulas in $\mathcal{L}_C^0 \cup \mathcal{L}^{fo}$ can be written as $\Delta_\rightarrow \cup \Delta_{fo}$, where $\Delta_\rightarrow \subseteq \mathcal{L}_C^0$ and $\Delta_{fo} \subseteq \mathcal{L}^{fo}$. Consider the following strengthening of LLE:

LLE$^+$. If $\vdash_{\Lambda \cup \Delta_{fo}} \varphi \Leftrightarrow \psi$, then from $\Delta \hookrightarrow \varphi_1 \rightarrow \psi$ infer $\Delta \hookrightarrow \varphi_2 \rightarrow \psi$.

RW can be similarly strengthened to RW$^+$.

Some rules from AX_C^Λ to deal with the universal quantification are also needed, specifically, variants of Λ-AX, F1, and F3, and another rule similar in spirit to F3:

Λ-AX$^+$. If $\vdash_{\Lambda \cup \Delta_{fo}} \varphi$, then $\Delta \hookrightarrow \varphi$.

F1$^+$. From $\forall x \varphi$ infer $\varphi[x/z]$, where z is a variable that does not appear in φ.

F3$^+$. If x does not appear free in Δ, then from $\Delta \hookrightarrow \varphi$ infer $\Delta \hookrightarrow \forall x \varphi$.

EQ. If x does not appear free in Δ, φ, or ψ, and σ is a first-order formula, then from $\Delta \cup \{\sigma\} \hookrightarrow \varphi$ infer $\Delta \cup \{\exists x \sigma\} \hookrightarrow \varphi$ (existential quantification).

[1] Λ is not usually mentioned explicitly, but it will be useful to do so for the results of this paper.

REN. If $y_1, \ldots, y_n$ do not appear in φ, then from $\forall x_1, \ldots, x_n \varphi$ infer $\forall y_1, \ldots, y_n(\varphi[x_1/y_1, \ldots, x_n/y_n])$ (renaming).

But these rules do not seem to suffice. Intuitively, what is needed is a way to capture the fact that the domain is the same in all worlds. In AX_C^Λ, the one axiom that captures this is F5. Unfortunately, F5 is not expressible in $\mathcal{L}_C^0$. To capture its effects in $\mathcal{L}_C^0$, a somewhat more complicated rule seems necessary.

Definition 3.4: An *interpretation-independent* formula φ is a first-order formula that does not mention any constant, function, or predicate symbols (and, thus, is a formula whose atomic predicates all are of the form $x = y$).

The following rule can be viewed as a variant of the OR rule for interpretation-independent formulas.

II. If $\Delta \cup \{\sigma_1\} \hookrightarrow \varphi$, $\Delta \cup \{\sigma_2\} \hookrightarrow \varphi$, and σ_1 and σ_2 are interpretation-independent, then $\Delta \cup \{\sigma_1 \vee \sigma_2\} \hookrightarrow \varphi$ (interpretation independence).

Let $\mathbf{P}_\Lambda^+$ consist of $\mathbf{P}_\Lambda$ (with LLE and RW replaced by LLE$^+$ and RW$^+$, respectively) together with F1$^+$, F3$^+$, EQ, REN, and II.

Theorem 3.5: *If $\Delta \cup \{\varphi\} \subseteq \mathcal{L}_C^0 \cup \mathcal{L}^{fo}$, then the following are equivalent:*

(a) $\mathbf{P}_\Lambda^+ \vdash \Delta \hookrightarrow \varphi$;

(b) $\mathcal{PS}(\Lambda) \models \Delta \hookrightarrow \varphi$;

(c) $\mathcal{PS}(\Lambda) \models^{sp} \Delta \hookrightarrow \varphi$.

Proof: The argument for soundness (that is, that (a) implies (c)) for the axioms and rules that also appear in $\mathbf{P}_\Lambda$ is essentially done in the proof of Theorem 3.2; the soundness of F1$^+$, F3$^+$, EQ, and REN is straightforward. The soundness of II follows easily from the observation that, since there is a fixed domain, if σ_1 and σ_2 are interpretation independent and $(M, V) \models \sigma_2 \vee \sigma_2$, then $(M, V) \models \sigma_1$ or $(M, V) \models \sigma_2$. This would not be the case for a formula such as $\mathbf{d}_1 = \mathbf{d}_2 \vee \mathbf{d}_1 = \mathbf{d}_3$. It could be that, for every world w, $(M, V, w) \models \mathbf{d}_1 = \mathbf{d}_2 \vee \mathbf{d}_1 = \mathbf{d}_3$, with either $\mathbf{d}_1 = \mathbf{d}_2$ being true in every world or $\mathbf{d}_1 = \mathbf{d}_3$ being true in every world.

The fact that (c) implies (b) follows just as in the proof of Theorem 3.2, using Theorem 3.1. Thus, it remains to show that (b) implies (a). As usual, for completeness, it suffices to show that if $\mathbf{P}_\Lambda^+ \not\vdash \Delta \hookrightarrow \forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$, then there is a structure $M \in \mathcal{PS}(\Lambda)$ and valuation V such that $M \models \Delta$ and $(M, V) \models \neg\forall x_1 \ldots \forall x_n(\varphi \rightarrow \psi)$. The idea is to reduce to Theorem 3.3 by transforming to a situation where the rules in $\mathbf{P}_\Lambda$ suffice. Let $Dist_k$ be the formula that says that there are at least k distinct domain elements;

$$\exists x_1 \ldots \exists x_k(\wedge_{1 \leq i < j \leq k} x_i \neq x_j).$$

Note that $Dist_k$ is interpretation-independent. Let $\Delta_0 = \Delta$; let $\Delta_{n+1} = \Delta_n \cup \{Dist_{n+1}\}$ if $\mathbf{P}_\Lambda^+ \not\vdash \Delta_n \cup \{Dist_{n+1}\} \hookrightarrow \varphi \rightarrow \psi$ and let $\Delta_{n+1} = \Delta_n \cup \{\neg Dist_{n+1}\}$ otherwise; finally, let $\Delta_\infty = \cup_n \Delta_n$.

An easy induction using II and Λ-AX$^+$ shows that $\mathbf{P}_\Lambda^+ \not\vdash \Delta_n \hookrightarrow \forall x_1 \ldots \forall x_n(\varphi^* \rightarrow \psi^*)$ for all n, and hence $\mathbf{P}_\Lambda^+ \not\vdash \Delta_\infty \hookrightarrow \forall x_1 \ldots \forall x_n(\varphi^* \rightarrow \psi^*)$. Let k^* be the largest k such that $Dist_k \in \Delta_\infty$ (where take $k^* = \infty$ if, for all k, $Dist_k \in \Delta_\infty$). Intuitively, k^* will be the size of the domain in the PS structure that we construct.

By REN, I can assume without loss of generality that $x_1, \ldots, x_n$ do not appear in Δ_∞. Thus, from F3$^+$, it follows that $\mathbf{P}_\Lambda^+ \not\vdash \Delta_\infty \hookrightarrow \varphi \rightarrow \psi$. A formula is a *(complete) equality statement* for $x_1, \ldots, x_n$ if it is a conjunction of formulas of the form $x_i = x_j$ and $x_i \neq x_j$, such that for all $1 \leq i < j \leq n$, either $x_i = x_j$ or $x_i \neq x_j$ is a conjunct. Note that a complete equality statement is interpretation-independent, and some equality statement must be true of every valuation. It thus follows (using II and Λ-AX$^+$), that $\mathbf{P}_\Lambda^+ \not\vdash \Delta_\infty \cup \{\sigma\} \hookrightarrow \varphi \rightarrow \psi$, for some equality statement σ. An equality statement σ partitions $x_1, \ldots, x_n$ into equivalence classes, where it follows from σ that all the variables in each equivalence class are equal to each other, but variables in two different equivalence classes are not equal to each other. Suppose that there are h equivalence classes. Clearly $h \leq k^*$ (for otherwise σ would be inconsistent with $dist_{k^*} \in \Delta_\infty$, so it would follow that $\mathbf{P}_\Lambda^+ \vdash \Delta_\infty \cup \{\sigma\} \hookrightarrow \varphi \rightarrow \psi$). Without loss of generality, I can assume that $x_1, \ldots, x_h$ are in distinct equivalence classes (so that σ implies that, for all j, $x_1 = x_j \vee \ldots \vee x_h = x_j$). Let φ' and ψ' be the result of replacing x_j for $j > h$ by x_i for $i \leq h$, where $\sigma \Rightarrow x_i = x_j$. It easily follows, using LLE$^+$ and RW$^+$, that $\mathbf{P}_\Lambda^+ \not\vdash \Delta_\infty \cup \{\sigma\} \hookrightarrow \varphi' \rightarrow \psi'$. Let y_i, $1 \leq i < k^* + 1 - h$ be fresh variables (where $k^* + 1 - h = \infty$ if $k^* = \infty$) that do not appear in Δ, φ', or ψ'. Let $\Delta_\infty^0 = \Delta_\infty \cup \{\sigma\}$; if $m > 0$, let $\Delta_\infty^m = \Delta_\infty^{m-1} \cup \{(\wedge_{1 \leq j < m} y_m \neq y_j \wedge (\wedge_{1 \leq j \leq h} y_m \neq x_j)\}$. Let $\Delta^\dagger = \cup_{i < k^*+1-h} \Delta_\infty^i$. I claim that $\mathbf{P}_\Lambda^+ \not\vdash \Delta^\dagger \hookrightarrow \varphi' \rightarrow \psi'$. To show this, since proofs are finite, it suffices to show that $\mathbf{P}_\Lambda^+ \not\vdash \Delta_\infty^m \hookrightarrow \varphi' \rightarrow \psi'$ for all $m < k^* + 1 - h$. I do this by induction on m. For $m = 0$, this is true by assumption. Suppose, by way of contradiction, that $\mathbf{P}_\Lambda^+ \vdash \Delta_\infty^m \hookrightarrow \varphi' \rightarrow \psi'$. For $m > 0$, note that Δ_∞^m has the form $\Delta_\infty \cup \{\sigma^0, \sigma^1, \ldots, \sigma^m\}$, where y_m appears only in σ^m. It thus follows from EQ that $\mathbf{P}_\Lambda^+ \vdash \Delta_\infty^m \cup \{\sigma^0, \ldots, \sigma^{m-1}, \exists y_m \sigma^m\} \hookrightarrow \varphi' \rightarrow \psi'$. It is easy to see that $Dist_m \Rightarrow \exists y_m \sigma^m$ is valid. Since $Dist_m \in \Delta_\infty$, it follows that $\mathbf{P}_\Lambda^+ \vdash \Delta_\infty^{m-1} \hookrightarrow \varphi' \rightarrow \psi'$, contradicting the inductive hypothesis.

Let $\mathbf{D} = \{\mathbf{d}_i : 1 \leq i < k^* + 1\}$ consist of fresh constant symbols not in $\mathcal{T} \cup \{\mathbf{c}_1, \ldots, \mathbf{c}_n\}$. Let an *instantiation* of $\forall x_1 \ldots \forall x_n(\varphi' \rightarrow \psi')$ be a formula in $\mathcal{L}_C^-$ of the form $\varphi'' \rightarrow \psi''$ that results by replacing each free variable in $\varphi' \rightarrow \psi'$ by some element of $\mathbf{D}$. For example, the instantiations of $\forall x \forall y(P(x, y) \rightarrow Q(y))$ are all formulas the form $P(\mathbf{d}_i, \mathbf{d}_j) \rightarrow Q(\mathbf{d}_j)$. Let Δ^* be the result of replacing each formula $\forall x_1 \ldots \forall x_n(\varphi' \rightarrow \psi')$ in Δ by all its instantiations, and replacing the free variables $x_1, \ldots, x_h, y_1, y_2, \ldots$ in the formulas in $\Delta^\dagger - \Delta$ by $\mathbf{d}_1, \mathbf{d}_2, \ldots$ respectively. Let φ^* and ψ^* be the result of replacing $x_1, \ldots, x_h$ in φ' and ψ' by $\mathbf{d}_1, \ldots, \mathbf{d}_h$, respectively. I claim that $\mathbf{P}_{\Lambda \cup \Delta_0^*} \not\vdash \Delta_\rightarrow^* \hookrightarrow \varphi^* \rightarrow \psi^*$. Suppose, by way of contradiction, that $\mathbf{P}_{\Lambda \cup \Delta_0^*} \vdash \Delta_\rightarrow^* \hookrightarrow \varphi^* \rightarrow \psi^*$. Then clearly $\mathbf{P}_\Lambda^+ \vdash \Delta^* \hookrightarrow \varphi^* \rightarrow \psi^*$. Let Δ_1 be the result of replacing all occurrences of $\mathbf{d}_1, \mathbf{d}_2, \ldots$ by $x_1, \ldots, x_h, y_1, y_2, \ldots$, respectively. Then $\mathbf{P}_\Lambda^+ \vdash \Delta_1 \hookrightarrow$

$\varphi' \rightarrow \psi'$ (simply replace the constants $\mathbf{d}_j$ by the appropriate variables in each line of the derivation of $\varphi^* \rightarrow \psi^*$). Moreover, Δ_1 has the form $\Delta_2 \cup (\Delta^\dagger - \Delta)$, and, by F1$^+$ and Λ-AX$^+$, it easily follows that $\Delta \cup (\Delta^\dagger - \Delta) \hookrightarrow \varphi' \rightarrow \psi'$; that is, $\Delta^\dagger \hookrightarrow \varphi' \rightarrow \psi'$. This gives us the desired contradiction.

Since $\mathbf{P}_{\Lambda \cup \Delta_0^*} \not\vdash \Delta^*_{\rightarrow} \hookrightarrow \varphi^* \rightarrow \psi^*$, it follows by Theorem 3.3 that there exists a PS structure $M \in \mathcal{PS}(\Lambda)$ and a valuation V such that $M \models \Delta^*$ and $M \models \varphi^* \rightarrow \psi^*$. Moreover, by standard arguments (used, for example, in (Friedman, Halpern, & Koller 2000)) the domain of M can be taken to be countable. Since $Dist_k \in \Delta^*$ for all $k < k^*$, and if $k^* < \infty$, $\neg Dist_{k^*} \in \Delta^*$, all the worlds in M mus+t have domain size k^* (where, if $k^* = \infty$, the domain is countable). Hence, I can assume without loss of generality that all the worlds have the same domain, which can be take to be $D = \{d_i : i < k^*\}$; moreover, d_i can be taken to be the interpretation $\mathbf{d}_i$ in each domain. It follows that $M \models \Delta$, and $M \not\models \varphi \rightarrow \psi$, as desired. ∎

3.2 Quantitative Reasoning

The super-polynomial semantics just talks about asymptotic complexity. It says that for all k, the conclusion will hold with probability greater than $1 - 1/n^k$ for sufficiently large n, provided that the assumptions hold with sufficiently high probability, where n can be, for example, the security parameter. While this asymptotic complexity certainly gives insight into the security of a protocol, in practice, a system designer wants to achieve a certain level of security, and needs to know, for example, how large to take the keys in order to achieve this. In this section, I provide a more quantitative semantics appropriate for such reasoning, and connect the qualitative and quantitative semantics.

The syntax of the quantitative language, which is denoted $\mathcal{L}_{C,q}$, is just like that of the qualitative language, except that, instead of formulas of the form $\varphi \rightarrow \psi$, there are formulas of the form $\varphi \rightarrow^r \psi$, where r is a real number in $[0, 1]$. The semantics of such a formula is straightforward:

$(M, V) \models \varphi \rightarrow^r \psi$ if there exists some $n^* \geq 0$ such that for all $n \geq n^*$, $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - r)$.

Let $\mathcal{L}^0_{C,q}$ be the obvious analogue of $\mathcal{L}^0_C$, consisting of all formulas of the form $\forall x_1 \ldots \forall x_n(\varphi \rightarrow^r \psi)$, where φ and ψ are first-order formulas.

For each of the axioms and rules in system $\mathbf{P}_\Lambda$, there is a corresponding sound axiom or rule in $\mathcal{L}^0_{C,q}$:

LLE^q. If $\vdash_{\Lambda \cup \Delta_{fo}} \varphi_1 \Leftrightarrow \varphi_2$, then from $\Delta \hookrightarrow \varphi_1 \rightarrow^r \psi$ infer $\Delta \hookrightarrow \varphi_2 \rightarrow^r \psi$.

RW^q. If $\vdash_{\Lambda \cup \Delta_{fo}} \psi_1 \Rightarrow \psi_2$, then from $\Delta \hookrightarrow \varphi \rightarrow^r \psi_1$ infer $\Delta \hookrightarrow \varphi \rightarrow^r \psi_2$.

REF^q. $\varphi \rightarrow^0 \varphi$ (reflexivity).

AND^q. From $\varphi \rightarrow^{r_1} \psi_1$ and $\varphi \rightarrow^{r_2} \psi_2$ infer $\varphi \rightarrow^{r_3} \psi_1 \wedge \psi_2$, where $r_3 = \min(r_1 + r_2, 1)$.

OR^q. From $\varphi_1 \rightarrow^{r_1} \psi$ and $\varphi_1 \rightarrow^{r_2} \psi$ infer $\varphi_1 \vee \varphi_2 \rightarrow^{r_3} \psi$, where $r_3 = \min(\max(2r_1, 2r_2), 1)$.

CM^q. From $\varphi_1 \rightarrow^{r_1} \varphi_2$ and $\varphi_1 \rightarrow^{r_2} \psi$ infer $\varphi \wedge \varphi_2 \rightarrow^{r_3} \psi$, where $r_3 = \min(r_1 + r_2, 1)$.

Let $\mathbf{P}^{+,q}_\Lambda$ denote this set of rules, together with F1$^+$, F3$^+$, EQ, REN, and II (all of which hold with no change in the quantitative setting), and

INC. If $r_1 \leq r_2$, then from $\varphi \rightarrow^{r_1} \psi$ infer $\varphi \rightarrow^{r_2} \psi$.

Theorem 3.6: *The rules in $\mathbf{P}^{+,q}_\Lambda$ are all sound.*

Proof: The soundness of the quantitative analogues of the rules in $\mathbf{P}_\Lambda$ is immediate from the proof of Theorem 3.2. The soundness of remaining rules holds as it did before (since they are unchanged). ∎

I do not believe that $\mathbf{P}^{+,q}_\Lambda$ is complete, nor do I have a candidate complete axiomatization for the quantitative language. Nevertheless, as the proofs in (Datta *et al.* 2008) show, $\mathbf{P}^{+,q}_\Lambda$ suffices for proving many results of interest in security. Moreover, as I now show, there is a deep relationship between $\mathbf{P}^+_\Lambda$ and $\mathbf{P}^{+,q}_\Lambda$. To make it precise, given a set of formulas $\Delta \subseteq \mathcal{L}^0_C$, say that $\Delta' \subseteq \mathcal{L}^0_{C,q}$ is a *quantitative instantiation* of Δ if, for every formula $\varphi \rightarrow \psi \in \Delta$, there is a bijection f from Δ to Δ' such that, for every formula $\varphi \rightarrow \psi \in \Delta$, there is a real number $r \in [0, 1]$ such that $f(\varphi \rightarrow \psi) = \varphi \rightarrow^r \psi$. That is, Δ' is a quantitative instantiation of Δ if each qualitative formula in Δ has a quantitative analogue in Δ'.

Although the proof of the following theorem is straightforward, it shows the power of using of $\mathbf{P}^+_\Lambda$. Specifically, it shows that if $\varphi \rightarrow \psi$ is derivable from Δ in $\mathbf{P}^+_\Lambda$ then, for all $r \in [0, 1]$, there exists a quantitative instantiation Δ' of Δ such that $\varphi \rightarrow^r \psi$ is derivable from Δ' in $\mathbf{P}^{+,q}_\Lambda$. Thus, if the system designer wants security at level r (that is, she wants to know that the desired security property holds with probability at least $1 - r$), then if she has a qualitative proof of the result, she can compute the strength with which her assumptions must hold in order for the desired conclusion to hold. For example, she can compute how to set the security parameters in order to get the desired level of security. This result can be viewed as justifying qualitative reasoning. Roughly speaking, it says that it is safe to avoid thinking about the quantitative details, since they can always be derived later. Note that this result would not hold if the language allowed negation. For example, even if $\neg(\varphi \rightarrow \psi)$ could be proved given some assumptions (using the axiom system AX$^\Lambda_C$), it would not necessarily follow that $\neg(\varphi \rightarrow^q \psi)$ holds, even if the probability of the assumptions was taken arbitrarily close to one.

Theorem 3.7: *If $\mathbf{P}^+_\Lambda \vdash \Delta \hookrightarrow \varphi \rightarrow \psi$, then for all $r \in [0, 1]$, there exists a quantitative instantiation Δ' of Δ such that $\mathbf{P}^{+,q}_\Lambda \vdash \Delta' \hookrightarrow \varphi \rightarrow^q \psi$. Moreover, Δ' can be found in polynomial time, given the derivation of $\Delta \hookrightarrow \varphi \rightarrow \psi$.*

Proof: The existence of Δ' follows by a straightforward induction on the length of the derivation. If it has length 1, then either $\varphi = \psi$, in which case it is an instance of REF^q, or $\varphi \rightarrow \psi \in \Delta$, in which case I simply take Δ' such that it includes $\varphi \rightarrow^r \psi$. For the inductive step, the only nontrivial case is if $\varphi \rightarrow \psi$ follows from earlier steps by an instance of a rule of inference in $\mathbf{P}_\Lambda^+$. The result then follows by a simple case analysis on the form of the rule. For example, if the AND rule is used, then $\varphi{\rightarrow}\psi$ has the form $\varphi{\rightarrow}\psi_1 \wedge \psi_2$, and there are shorter derivations of $\Delta \hookrightarrow \varphi{\rightarrow}\psi_1$ and $\Delta \hookrightarrow \varphi{\rightarrow}\psi_2$. Choose $s_1, s_2 \in [0,1]$ such that $s_1+s_2 = r$.[2] By the induction hypothesis, there exist variants Δ_1 and Δ_2 such that $\mathbf{P}_\Lambda^{+,q} \vdash \Delta_i \hookrightarrow \varphi{\rightarrow}^{s_i}\psi_i$, for $i = 1,2$. Let Δ_3 be a quantitative instantiation of Δ that dominates both Δ_1 and Δ_2, in the sense that if $\varphi'{\rightarrow}^{r_i}\psi' \in \Delta_i$, for $i = 1,2,3$, then $r_i \geq \max(r_1, r_2)$. Then it is easy to see that $\mathbf{P}_\Lambda^{+,q} \vdash \Delta_3 \hookrightarrow \varphi{\rightarrow}^{s_i}\psi_i$ for $i = 1,2$. By AND^q, it easily follows that $\mathbf{P}_\Lambda^{+,q}\Delta_3 \hookrightarrow \varphi{\rightarrow}^r\psi_1 \wedge \psi_2$. The argument for all the other rules in $\mathbf{P}_\Lambda^+$ is similar. This argument also shows that finding Δ' from the proof of $\Delta \hookrightarrow \varphi{\rightarrow}\psi$ just involves solving some simple linear inequalities, which can be done in polynomial time. ∎

The proof of Theorem 3.7 gives even more useful information to the system designer. In general, there may be a number of quantitative instantiations Δ' of Δ that give the desired conclusion. For example, as the proof shows, if the AND rule is used in the qualitative proof, and we want the conclusion to hold at level r, we must just choose s_1 and s_2 such that $\varphi{\rightarrow}\psi_1$ and $\varphi{\rightarrow}\psi_2$ hold at level s_1 and s_2, respectively. If the system designer finds it easier to satisfy the first formula than the second (for example, the first may involve the length of the key, while the second may involve the degree of trustworthiness of one of the participants in the protocol), there may be an advantage in choosing s_1 relatively small and s_2 larger. As long as $s_1 + s_2 = r$, the desired conclusion will hold.

Acknowledgments: I think Anupam Datta, John Mitchell, Riccardo Pucella, and Arnab Roy for many useful discussions on applying conditional logic to security protocols.

References

Abadi, M., and Rogaway, P. 2000. Reconciling two views of cryptography (the computational soundness of formal encryption). In *Proc. IFIP International Conference on Theoretical Computer Science (TCS'00)*, volume 1872 of *Lecture Notes in Computer Science*, 3–22. Springer-Verlag.

Adams, E. 1975. *The Logic of Conditionals*. Dordrecht, Netherlands: Reidel.

Bella, G., and Paulson, L. C. 1998. Kerberos version IV: Inductive analysis of the secrecy goals. In Quisquater, J.-J., ed., *Proc. 5th European Symposium on Research in Computer Security*, LNCS, Volume 1485, 361–375. Springer-Verlag.

Bellare, M.; Canetti, R.; and Krawczyk, H. 1998. A modular approach to the design and analysis of authentication and key exchange protocols. In *Proc. 30th Annual Symposium on the Theory of Computing*.

Borisov, N.; Goldberg, I.; and Wagner, D. 2001. Intercepting mobile communications: the insecurity of 802.11. In *Proc. 7th Annual International Conference on Mobile Computing and Networking*, 180–189.

Burgess, J. 1981. Quick completeness proofs for some logics of conditionals. *Notre Dame Journal of Formal Logic* 22:76–84.

Datta, A.; Derek, A.; Mitchell, J. C.; Shmatikov, V.; and Turuani, M. 2005. Probabilistic polynomial-time semantics for a protocol security logic. In *32nd International Colloquium on Automata, Languages, and Programming (ICALP)*, 16–29.

Datta, A.; Derek, A.; Mitchell, J. C.; and Roy, A. 2007. Protocol composition logic (PCL). *Electronic Notes Theoretical Computer Science* 172:311–358.

Datta, A.; Halpern, J. Y.; Mitchell, J. C.; Pucella, R.; and Roy, A. 2008. Reasoning about conditional probability and concrete security in protocol proofs. Unpublished manuscript.

Enderton, H. B. 1972. *A Mathematical Introduction to Logic*. New York: Academic Press.

Friedman, N., and Halpern, J. Y. 2001. Plausibility measures and default reasoning. *Journal of the ACM* 48(4):648–685.

Friedman, N.; Halpern, J. Y.; and Koller, D. 2000. First-order conditional logic for default reasoning revisited. *ACM Trans. on Computational Logic* 1(2):175–207.

Geffner, H. 1992. High probabilities, model preference and default arguments. *Mind and Machines* 2:51–70.

Goldreich, O. 2001. *Foundations of Cryptography, Vol. 1*. Cambridge University Press.

Goldszmidt, M., and Pearl, J. 1992. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Principles of Knowledge Representation and Reasoning: Proc. Third International Conference (KR '92)*. 661–672.

Goldszmidt, M.; Morris, P.; and Pearl, J. 1993. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions of Pattern Analysis and Machine Intelligence* 15(3):220–232.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207.

Mitchell, J.; Mitchell, M.; and Stern, U. 1997. Automated analysis of cryptographic protocols using Murφ. In *Proc. 1997 IEEE Symposium on Security and Privacy*, 141–151. IEEE Computer Society Press.

Mitchell, J. C.; Shmatikov, V.; and Stern, U. 1998. Finite-state analysis of SSL 3.0. In *Proc. Seventh USENIX Security Symposium*, 201–216.

Paulson, L. C. 1994. *Isabelle, A Generic Theorem Prover*, volume 828 of *Lecture Notes in Computer Science*. Springer-Verlag.

Wagner, D., and Schneier, B. 1996. Analysis of the SSL 3.0 protocol. In *Proc. 2nd USENIX Workshop on Electronic Commerce*.

[2]It suffices to take $s_1 = s_2 = r/2$, but there is an advantage to having this greater flexibility; see the discussion after the proof.