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 and 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, and Wagner 2001] to standards and proposed standards for Secure Socket Layer [Wagner and Schneier 1996; Mitchell, Shmatikov, and Stern 1998] to Kerberos [Bella and 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

*A preliminary version of this paper appeared in *AAAI-08 (Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence)*, 2008. This version uses a somewhat different axiomatization than the conference version, and includes proofs not contained in the conference version, as well as more detailed discussion.

†Work supported in part by NSF under under grants IIS-0812045, IIS-0911036, and CCF-1214844 ITR-0325453 and IIS-0534064, and by AFOSR under grant FA9550-05-1-0055.

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, and 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, and 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, Derek, Mitchell, and Roy 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, and Magidor 1990]. Indeed, one semantics given to $\rightarrow$, called ϵ-*semantics* [Adams 1975; Goldszmidt and Pearl 1992], is very close in spirit to that used in [Datta, Derek, Mitchell, Shmatikov, and Turuani 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, $\lim_{n\rightarrow\infty} \Pr_n(\psi \mid \varphi) = 1$ (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 $\lim_{n\rightarrow\infty} \Pr_n(\psi \mid \varphi) = 1$.) In a companion paper [Datta, Halpern, Mitchell, Roy, and Sen 2015], it is shown that reinterpreting $\rightarrow$ in this way gives an elegant, powerful variant of the logic considered in [Datta, Derek, Mitchell, Shmatikov, and Turuani 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, and 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, Derek, Mitchell, Shmatikov, and Turuani 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 that 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 about 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. More generally, the proof shows that, given a qualitative proof of a property and ϵ, we can compute a δ in polynomial time from the proof itself such that if the assumptions in the proof all hold with probability at least $1 - \delta$, then the conclusion holds with probability at least $1 - \epsilon$. 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.

There has been work on formal proof techniques for concrete cryptography [Blanchet 2006; Barthe, Grégoire, and Zanella-Béguelin 2009; Barthe, Grégoire, Heraud, and Zanella-Béguelin 2011]. However, it has largely focused on relational techniques where security is proved via game-based reductions, similar to traditional cryptographic proofs. The kind of transition from qualitative to quantitative reasoning that is my focus here has not been investigated. Perhaps even closer to this paper is that of Bana, Hasebe, and Okada [2013]. They also have an operator $\rightarrow$ that can be viewed as attempting to capture qualitatively some probabilistic aspects of reasoning about security. I discuss their work in more detail in Section 4.

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. I conclude in Section 4 with some discussion of the usefulness of this logic for reasoning about security.

2 First-Order Conditional Logic

I review the syntax and semantics of 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 (i.e., closed quantifier-free first-order formulas) 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 $\mathcal{L}^0_C(\mathcal{T})$ 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 $\mathcal{L}^0_C$ does not include negations of $\rightarrow$ formulas or conjunctions of $\rightarrow$ formulas. While not having conjunctions does not really impair the expressive power of $\mathcal{L}^0_C$ (since we will be interested in sets of $\mathcal{L}^0_C$ 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) 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, where each probability measure $\Pr_n$ in the sequence is a probability measure on W. I assume for ease of exposition that all subsets of W are

measurable with respect to each probability measure $\Pr_n$ in the sequence. (I discuss below how this assumption can be weakened considerably.) As usual, a *valuation* V associates with each variable x an element $V(x) \in D$.

Given a valuation V and structure $M = (D, W, \pi, \mathcal{P})$, 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 $w \in W$, written $(M, V, w) \models \varphi$, is determined as usual. For $\varphi \in \mathcal{L}_C$, let $[\![\varphi]\!]_{M,V} = \{w \in W : (M, V, w) \models \varphi\}$. 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 and valuation V. Define

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

where $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V})$ is taken to be 1 if $\Pr_n([\![\varphi]\!]_{M,V}) = 0$. (I could have written $(M, V) \models \varphi \rightarrow \psi$, since the truth of $\varphi \rightarrow \psi$ is independent of the world w; I occasionally do this below.)

If W is infinite, the assumption that every subset of W is measurable is a very strong one. Suppose instead that there is a fixed σ-algebra $\mathcal{F}$ of measurable sets that is the domain of all the probability measures $\Pr_n$ in the sequence. All that is needed in the semantics above is that $[\![\varphi]\!]_{M,V}$ is measurable (i.e., $[\![\varphi]\!]_{M,V} \in \mathcal{F}$) for every formula φ and valuation V. I now give a condition sufficient to guarantee this.

As usual, the set of terms over the vocabulary $\mathcal{T}$ is defined inductively. A variable x is a term, a constant $c \in \mathcal{T}$ is a term, and if $f \in \mathcal{T}$ is a k-ary function symbol and $t_1, \ldots, t_k$ are terms, then $f(t_1, \ldots, t_k)$ is a term. An *atomic expression* over $\mathcal{T}$ has the form either (a) $P(t_1, \ldots, t_k)$, where $P \in \mathcal{T}$ is a k-ary predicate symbol and $t_1, \ldots, t_k$ are terms over $\mathcal{T}$, (b) $f(t_1, \ldots, t_k) = t_{k+1}$, where $f \in \mathcal{T}$ is a k-ary predicate symbol and $t_1, \ldots, t_{k+1}$ are terms over $\mathcal{T}$, or (c) $t_1 = t_2$, where t_1 and t_2 are terms over $\mathcal{T}$. I claim that if the domain D is finite or countably infinite and for all valuations V and atomic expressions A, $[\![A]\!]_{M,V}$ is measurable, then $[\![\varphi]\!]_{M,V}$ is measurable for all formulas φ in the language and valuations V. The claim follows by a straightforward induction on the structure of formulas. For a formula φ of the form $\forall x \varphi'$, observe that $[\![\forall x \varphi']\!]_{M,V} = \cap_{d \in D} [\![\varphi']\!]_{M,V_d}$, where V_d is the valuation that agrees with V except that $V_d(x) = d$. Since D is countable and $\mathcal{F}$ is closed under countable intersection, it follows that $[\![\forall x \varphi']\!]_{M,V} \in \mathcal{F}$. If φ has the form $\varphi' \rightarrow \psi$, note that $[\![\varphi' \rightarrow \psi]\!]_{M,V}$ is either $\emptyset$ or W, since it is independent of the world, so $[\![\varphi' \rightarrow \psi]\!]$ must be measurable. All other cases in the induction argument are straightforward.

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

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

Note that if $(M, V, w) \models^{sp} \varphi \rightarrow \psi$ and $p(n)$ is a polynomial whose leading coefficient is positive, then $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - \frac{1}{p(n)}$ 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 there exists n^* such that, for all $n \geq n^*$, $\Pr_n([\![\mathit{false}]\!]_M \mid [\![\neg\varphi]\!]_M) \geq 1/2$. Since $[\![\mathit{false}]\!]_M = \emptyset$, this can happen only if $\Pr([\![\neg\varphi]\!]_M) = 0$ for all $n \geq n^*$, or equivalently, $\Pr_n([\![\varphi]\!]_M) = 1$ 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 eventually 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. Formally, Λ is just a (possibly infinite) set of first-order formulas.

Let $\vdash_\Lambda$ denote provability in first-order logic given the formulas 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, and 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, and 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.

I now briefly discuss the axioms. Λ-AX just says that any first-order formula provable in Λ ia also provable in AX_C^{Λ}. The notion of "substitution instance of a propositional tautology" in C0 means that we start with a propositional tautology, and then uniformly replace each instance of a propositional variable by a formula in $\mathcal{L}_C(\mathcal{T})$. For example, $\forall x, y(P(x) \rightarrow Q(y)) \vee \neg \forall x, y(P(x) \rightarrow Q(y))$ is a substitution instance of the propositional tautology $p \vee \neg p$, where $\forall x, y(P(x) \rightarrow Q(y))$ is substituted for p.

C1 has been called *reflexivity*; it says, for example, that birds are typically birds. C2 is called the *and rule*; it says, for example, that if birds typically fly and birds typically have wings, then birds typically both fly and have wings. C3 is the *or rule*; it says, for example, that if birds typically fly and insects typically fly, that something that is either a bird or an insect typically flies. C4 is *cautious monotonicity*; it says that if birds typically fly and birds typically have wings, then flying birds typically have wings. Full monotonicity would say that if birds typically have wings then red birds have wings or, more generally, $(\varphi \rightarrow \psi) \Rightarrow ((\varphi \wedge \varphi') \rightarrow \psi)$, for an arbitrary φ'. This is not sound (although the analogue is sound if we replace $\rightarrow$ by $\Rightarrow$). R1 is *left logical equivalence*; it says that if we replace the left-hand side of a $\rightarrow$ formula by a provably equivalent formula, the resulting $\rightarrow$ formula is equivalent to the original formula. Again, the stronger version (that if we replace the left-hand side by a stronger formula we get a formula weaker than the original) does not hold; this is just full monotonicity. R2 is *right weakening*; it says that if we replace the right-hand side of a $\rightarrow$ formula by a weaker formula, the resulting $\rightarrow$ formula is weaker than the original formula. The axioms and rules C1–C4, R1, and R2 characterize the core properties of $\rightarrow$, and will essentially reappear in system **P**, discussed below.

C5 encodes the fact that the truth of $\varphi \rightarrow \psi$ depends only on the world; if it is true at some world, it is true at all worlds (and hence holds with probability 1). This lets us handle nested occurrence of $\rightarrow$. C6 is clearly a trivial consequence of the probabilistic semantics.

To understand the notion of "substitutable" in F1, observe that 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 $M = (D, W, \pi, \mathcal{P})$, where $W = \{w_1, w_2\}$, $D = \{d_1, d_2\}$, $\mathcal{P} = (\Pr_1, \Pr_2, \ldots)$ is such that $\Pr_n$ gives positive probability to both w_1 and w_2, and π is such 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 $NP(x)$ does not hold, no matter how x is interpreted. It follows that $M \models NP(c)$ and $M \models \forall x(\neg NP(x))$. 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.

The fact that variables are rigid is what makes F5 sound: if $x \neq y$ in some world given some valuation V, then $x \neq y$ in all worlds given valuation V, and hence holds with probability 1. F2, F3, F4, MP, and Gen are standard axioms and rules of first-order logic, and apply to modal logic as well.

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 such that 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 $\mathcal{P} = (\Pr_1, \Pr_2, \ldots)$. 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.) Define a valuation V to be *constant at k* if $V(x_m) = d_1$ for all $m \geq k$; a valuation V is *eventually constant* if V is constant at k for some k. Clearly there are only countably many eventually constant valuations. Let L^+ be an enumeration of pairs of the form $(\varphi \rightarrow \psi, V)$ such that V is eventually constant and $(M, V) \models \varphi \rightarrow \psi$; let $L^- = ((\varphi_1 \rightarrow \psi_1, V_1), (\varphi_2 \rightarrow \psi_2, V_2), \ldots)$ be an enumeration of pairs of the form $(\varphi \rightarrow \psi, V)$ such that V is eventually constant and $(M, V) \models \neg(\varphi \rightarrow \psi)$, where each pair appears in the enumeration L^- infinitely often. There exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for all n, if $n' > f(n)$, then $\Pr_{n'}(\llbracket \psi \rrbracket_{M,V} \mid \llbracket \varphi \rrbracket_{M,V}) \geq 1 - 1/n^n$ for all the first n pairs $(\varphi \rightarrow \psi, V) \in L^+$. Similarly, there exists a function g that maps elements in the enumeration L^- to $\mathbb{N}$ such that if $(\varphi \rightarrow \psi, V)$ is in L^-, then $g(\varphi \rightarrow \psi, V)$ is the least integer k such that, for infinitely many indices h, we have $\Pr_h(\llbracket \psi \rrbracket_{M,V} \mid \llbracket \varphi \rrbracket_{M,V}) < 1 - 1/k$. (There must be such an integer k, since $\lim_{h \rightarrow \infty} \Pr_h(\llbracket \psi \rrbracket_{M,V} \mid \llbracket \varphi \rrbracket_{M,V}) \neq 1$.)

I now construct a subsequence $\mathcal{P}' = (\Pr'_1, \Pr'_2, \ldots)$ of $\mathcal{P}$ by taking $\Pr'_n = \Pr_N$, where, if $(\varphi \rightarrow \psi, V)$ is the nth element of L^-, then N is the least integer greater than $f(n)$ such that $\Pr'_N(\llbracket \psi \rrbracket_{M,V} \mid \llbracket \varphi \rrbracket_{M,V}) < 1 - 1/g(\varphi \rightarrow \psi, V)$. 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 $\varphi' \rightarrow \psi'$, suppose that the free variables in $\varphi' \rightarrow \psi'$ are contained in $\{x_1, \ldots, x_k\}$. Given V, let V' be an eventually constant valuation such that $V(x_i) = V'(x_i)$ for $i = 1, \ldots, k$. Clearly, for all distributions Pr on W, we have $\Pr([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) = \Pr([\![\psi]\!]_{M,V'} \mid [\![\varphi]\!]_{M,V'})$, and similarly with M replaced by M'. First suppose that $(M, V) \models \varphi' \rightarrow \psi'$. Thus, we must have $(M, V') \models \varphi' \rightarrow \psi'$, so $(\varphi' \rightarrow \psi', V')$ is in L^+. Suppose that $(\varphi' \rightarrow \psi', V')$ is the Nth pair in L^+. The construction guarantees that for all $n \geq N$, we have $\Pr'_n([\![\psi]\!]_{M,V'} \mid [\![\varphi]\!]_{M,V'}) \geq 1 - 1/n^n$. It follows that $(M', V') \models^{sp} \varphi' \rightarrow \psi'$, and thus $(M', V) \models^{sp} \varphi' \rightarrow \psi'$.

Next suppose that $(M, V) \models \neg(\varphi' \rightarrow \psi')$. Thus, $(M, V') \models \neg(\varphi' \rightarrow \psi')$, so the pair $(\varphi' \rightarrow \psi', V')$ appears infinitely often in the enumeration L^-. For each index h such that this pair appears in position h in L^-, by construction, we have that $\Pr'_h([\![\psi']\!]_{M',V'} \mid [\![\varphi']\!]_{M',V'}) < 1 - 1/g(\varphi' \rightarrow \psi', V')$. Hence, $(M', V') \models^{sp} \neg(\varphi' \rightarrow \psi')$, so $(M', V) \models^{sp} \neg(\varphi' \rightarrow \psi')$, as desired.[1] ∎

Let $\mathcal{PS}(\Lambda)$ consist of all PS structures M such that every world in M 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, and 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 valid, and that the rules of inference preserve validity. This is straightforward for all the axioms and rules other than C2, C3, C4, and C5. 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 $n \geq \max(n_1^*, n_2^*)$, $\Pr_n([\![\neg\psi_i]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \leq 1/2p(n)$. Thus, for $n \geq \max(n_1^*, n_2^*)$,

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

[1] Exactly the same argument shows that there exists a probability sequence $\mathcal{P}'$ such that, for all valuations V, $(M, V) \models \varphi$ iff $(M', V) \models^{ex} \varphi$, where $M' = (D, W, \pi, \mathcal{P}')$ and $\models^{ex}$ considers exponential convergence for $\rightarrow$; that is, $(M, V, w) \models^{ex} \varphi \rightarrow \psi$ if, for all c, there exists some $n^* \geq 0$ such that, for all $n \geq n^*$, $\Pr_n([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - \frac{1}{2^{cn}}$. I have focused on super-polynomial rather than exponential convergence here since that is what is considered in the security literature.

For C3, note that

$$\begin{aligned}
& \Pr(A \mid B_1 \cup B_2) \\
= \; & \Pr((A \cap B_1 \cup A \cap B_2) \mid B_1 \cup B_2) \\
= \; & \Pr(A \cap B_1 \mid B_1 \cup B_2) + \Pr(A \cap B_2 \mid B_1 \cup B_2) - \Pr(A \cap B_1 \cap B_2 \mid B_1 \cup B_2) \\
= \; & \Pr(A \mid B_1) \times \Pr(B_1 \mid B_1 \cup B_2) + \Pr(A \mid B_2) \times \Pr(B_2 \mid B_1 \cup B_2) - \Pr(A \cap B_1 \cap B_2 \mid B_1 \cup B_2).
\end{aligned} \tag{1}$$

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 (1) that if $n \geq \max(n_1^*, n_2^*)$, then

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

For C4, note that

$$\Pr(A_1 \mid A_2 \cap B) = \Pr(A_1 \cap A_2 \mid B) / \Pr(A_2 \mid B) \geq \Pr(A_2 \cap A_2 \mid B),$$

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

Finally, the validity of 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, and 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 $\mathcal{PS}(\Lambda) \not\models^{sp} \varphi$.[2] ∎

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.) More precisely, I characterize when one formula in $\mathcal{L}_C^0$ can be derived from other

[2] An easy extension of this argument shows that conditions (a), (b), and (c) of Theorem 3.2 are also equivalent to $\mathcal{PS}(\Lambda) \models^{ex} \varphi$. The argument that the axioms are sound for the $\models^{ex}$ semantics is similar in spirit to that above showing that they are sound for the $\models^{sp}$ semantics; this shows that $\mathrm{AX}_C^{\Lambda} \vdash \varphi$ implies $\mathcal{PS}(\Lambda) \models^{ex} \varphi$. The fact that $\mathcal{PS}(\Lambda) \models^{ex} \varphi$ is equivalent to $\mathcal{PS}(\Lambda) \models \varphi$ follows from the observation in the previous footnote that Theorem 3.1 also holds for the $\models^{ex}$ semantics.

formulas in $\mathcal{L}_C^0$, given a first-order theory. I first consider 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. I start by giving a sound and complete axiomatization for expressions of the form $\Delta \hookrightarrow \varphi$, where φ is a formula in $\mathcal{L}_C^-$ and Δ is a set of formulas in $\mathcal{L}_C^-$. I allow Δ to be infinite here (this seems more consistent with the usage in [Kraus, Lehmann, and Magidor 1990]), but all the results hold without change if Δ is restricted to being finite. (I make comments along the way about the changes needed if Δ is restricted to being finite.) $(M, V) \models \Delta \hookrightarrow \varphi$ if $(M, V) \models \varphi'$ for every formula $\varphi' \in \Delta$ implies that $(M, V) \models \varphi$. If Δ is finite, then $(M, V) \models \Delta \hookrightarrow \varphi$ iff $(M, V) \models (\wedge_{\varphi' \in \Delta} \varphi') \Rightarrow \varphi$. Thus, if Δ is finite, the $\Delta \hookrightarrow \varphi$ is expressible in $\mathcal{L}_C$; I allow the slightly greater generality of infinite sets to be able to relate the results of this paper to earlier work. I write $\mathcal{PS}(\Lambda) \models \Delta \hookrightarrow \varphi$ if $(M, V) \models \Delta \hookrightarrow \varphi$ for all PS structures M and valuations V.

Consider the following axioms:

LLE. $\{\varphi_1 \rightarrow \psi\} \hookrightarrow (\varphi_2 \rightarrow \psi)$ if $\vdash_\Lambda \varphi_1 \Leftrightarrow \varphi_2$ (left logical equivalence).

RW. $\{\varphi \rightarrow \psi_1\} \hookrightarrow (\varphi \rightarrow \psi_2)$ if $\vdash_\Lambda \psi_1 \Rightarrow \psi_2$ (right weakening).

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

AND. $\{\varphi \rightarrow \psi_1, \varphi \rightarrow \psi_2\} \hookrightarrow (\varphi \rightarrow (\psi_1 \wedge \psi_2))$.

OR. $\{\varphi_1 \rightarrow \psi, \varphi_2 \rightarrow \psi\} \hookrightarrow ((\varphi_1 \vee \varphi_2) \rightarrow \psi)$.

CM. $\{\varphi_1 \rightarrow \varphi_2, \varphi_1 \rightarrow \psi\} \hookrightarrow ((\varphi_1 \wedge \varphi_2) \rightarrow \psi)$ (cautious monotonicity).

TRIV. $\Delta \hookrightarrow \varphi$ if $\varphi \in \Delta$ (trivial)

We have one rule of inference:

TRANS. From $\Delta \hookrightarrow \psi$ for all $\psi \in \Delta'$ and $\Delta' \hookrightarrow \varphi$ infer $\Delta \hookrightarrow \varphi$ (transitivity).

This collection of rules has been called system $\mathbf{P}_\Lambda$ [Kraus, Lehmann, and Magidor 1990] or *the KLM properties*.[3] The rules are obvious analogues of axioms in AX_C^Λ. In particular, LLE is the analogue of R1; and RW is the analogue of R2; REF, AND, and OR are just restatements of C2, C3, and C4, respectively, in this notation;

Note that even if Δ and Λ are infinite, since all the rules are in fact finitary, we have $\vdash_\Lambda \Delta \hookrightarrow \varphi$ iff there is a finite set $\Lambda' \subseteq \Lambda$ and a finite set $\Delta' \subseteq \Delta$ such that $\vdash_{\Lambda'} \Delta' \hookrightarrow \varphi$. For the completeness result, I make an innocuous technical restriction: I assume that the vocabulary $\mathcal{T}$ includes a countably infinite set of constants and there

[3]The standard presentation of system $\mathbf{P}_\Lambda$ is somewhat different from that given here. For one thing, Λ is not usually mentioned explicitly; I mention it here to make the dependence on Λ clear. For another, the set Δ is typically not included as part of the rule, but is put on the left-hand side of $\vdash$, so what I am calling an axiom is typically viewed as an inference rule; for example, AND is usually written as "from $\varphi \rightarrow \psi_1$ and $\varphi \rightarrow \psi_2$ infer $\varphi \rightarrow (\psi_1 \wedge \psi_2)$". Finally, the axiom TRIV and the rule TRANS are not usually given explicitly, but are built into how inference works in $\mathbf{P}_\Lambda$. However, it is easy to see that what I have done here is equivalent to the standard approach.

are (as usual) countably many variables, and for every formula $\Delta \hookrightarrow \varphi$, there are infinitely many constants and variables that do *not* appear in Δ.[4]

The following result is well known.

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

I want to extend this result from $\mathcal{L}_C^-$ to $\mathcal{L}_C^0$; thus, I now want to allow $\Delta \cup \{\varphi\} \subseteq \mathcal{L}_C^0$. I actually extend a little further to allow first-order formulas, so $\Delta \cup \{\varphi\} \subseteq \mathcal{L}^{fo} \cup \mathcal{L}_C^0$. In addition, I want the axiomatization to be sound and complete for the $\models^{sp}$ semantics as well as the $\models$ semantics, so as to make it more applicable to reasoning about security protocols.

To get a complete axiomatization, I still need to use the axioms and rules of $\mathbf{P}_\Lambda$, but now I need (special cases of) a few other axioms and rules in AX_C^Λ modified to deal with this language:

C6$^+$. $\{true \rightarrow false\} \hookrightarrow false$.

F1$^+$. $\{\forall x \varphi\} \hookrightarrow \varphi[x/t]$ if t is substitutable for x in φ.

F4$^+$. $\{x = y, \varphi_1\} \hookrightarrow \varphi_2$, 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$^+$. $\{x \neq y\} \hookrightarrow N(x \neq y)$.

IMP. $\Delta \hookrightarrow \varphi$ if $\Delta \cup \{\varphi\} \subseteq \mathcal{L}^{fo}$, Δ is finite, and $\vdash_\Lambda \wedge_{\psi \in \Delta} \psi \Rightarrow \varphi$ (implication).

Gen$^+$. If z is a variable that does not appear free in Δ, then (a) if $\varphi \in \mathcal{L}^{fo}$, then from $\Delta \cup \{\varphi[x/z]\} \hookrightarrow \psi$ infer $\Delta \cup \{\exists x \varphi\} \hookrightarrow \psi$; (b) from $\Delta \hookrightarrow \varphi[x/z]$ infer $\Delta \hookrightarrow \forall x \varphi$.

Of course, C6$^+$, F4$^+$ and F5$^+$ are just C6, F4, and F5 restated using $\hookrightarrow$, F1$^+$ is just a special case of F1 (restated using $\hookrightarrow$), IMP is a special case of Λ-AX, and Gen$^+$ can be viewed as a special case of Gen. (It is not hard to show that it follows from Gen if Δ is finite.) Since I now allow first-order formulas, whose truth depends on the world, I take $(M, V, w) \models \Delta \hookrightarrow \varphi$ if $(M, V, w) \models \varphi'$ for all formulas $\varphi' \in \Delta$ implies $(M, V, w) \models \varphi$. I now write $\mathcal{PS}(\Lambda) \models \Delta \hookrightarrow \varphi$ if $(M, V, w) \models \Delta \hookrightarrow \varphi$ for all PS structures M and valuations V.

Let $\mathbf{P}_\Lambda^+$ be the axiom system consisting of the axioms and rules of $\mathbf{P}_\Lambda$ together with C6$^+$, F1$^+$, F4$^+$, F5$^+$, IMP, and Gen$^+$. I write $\mathbf{P}_\Lambda^+ \vdash \Delta \hookrightarrow \varphi$ if there is a derivation from $\mathbf{P}_\Lambda^+$ whose last line in $\Delta \hookrightarrow \varphi$.

Theorem 3.4: *If* $\Delta \cup \{\varphi\} \subseteq \mathcal{L}_C^0 \cup \mathcal{L}^{fo}$, *then the following are equivalent:*

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

[4]The assumption holds trivially if Δ is restricted to being a finite set, as long as the set of constants and variables is infinite. Even if Δ can be infinite, we can always add countably infinite fresh constants to $\mathcal{T}$ and rename variables, so this restriction is innocuous. The assumption is used in the proof of Theorem 3.4 below, where at one point I need to assume that there infinitely many "fresh" constants and variables that do not appear in Δ or φ. A similar assumption arises in the proof of Theorem 3.3 .

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

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

Proof: The fact that (a) implies (c) follows from the proof of Theorem 3.2, since all the axioms and rules in $\mathbf{P}^+_\Lambda$ are essentially (special cases of) axioms and rules in AX^Λ_C. 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 $\mathbf{P}^+_\Lambda \not\vdash \Delta \hookrightarrow \varphi$, then there is a structure $M \in \mathcal{PS}(\Lambda)$, a valuation V, and a world w such that $(M, V, w) \models \Delta$ and $(M, V, w) \models \neg\varphi$; roughly speaking, this says that if $\Delta \hookrightarrow \varphi$ is not provable in $\mathbf{P}^+_\Lambda$, then its negation is satisfiable.

The proof is quite similar in spirit to the completeness proof for AX^Λ_C given in [Friedman, Halpern, and Koller 2000], but there are some subtleties involved in dealing with the restricted language. Specifically, the proof in [Friedman, Halpern, and Koller 2000] uses a Henkin-style argument, using maximal consistent sets C of formulas. Because we work here with formulas of the form $\Delta \hookrightarrow \varphi$, we must redefine maximal consistent sets appropriately.

A set Δ' of formulas in $\mathcal{L}^{fo} \cup \mathcal{L}^0_C$ is $\mathbf{P}^+_\Lambda$*-consistent* if there is no finite subset Δ'' of Δ' such that $\vdash_\Lambda \Delta'' \hookrightarrow \mathit{false}$. If $\mathcal{T}^*$ is a vocabulary and $\mathcal{Y}$ is a set of constants and variables, then a *maximal* $\mathbf{P}^+_\Lambda$*-*$(\mathcal{T}^* \cup \mathcal{Y})$*–consistent* set of formulas is a $\mathbf{P}^+_\Lambda$-consistent set Δ' of formulas that use only symbols in $\mathcal{T}^* \cup \mathcal{Y}$ such that if ψ is a formula that uses only symbols in $\mathcal{T}^* \cup \mathcal{Y}$, then $\Delta' \cup \{\psi\}$ is not $\mathbf{P}^+_\Lambda$-consistent. Δ' is $\mathcal{T}^*$*-*$\mathcal{Y}$*–good* if (1) Δ' is a maximal Λ–$(\mathcal{T}^* \cup \mathcal{Y})$-consistent set of formulas, (2) $\exists x\psi \in \Delta' \cap \mathcal{L}^{fo}$ implies $\neg\psi[x/y] \in \Delta'$ for some $y \in \mathcal{Y}$, (3) if $\forall x\varphi \in \Delta'$ then $\varphi[x/y] \in \Delta'$ for all $y \in \mathcal{Y}$, and (4) if $y = y'$ (resp., $y \neq y'$) is in Δ for $y, y' \in \mathcal{Y}$, then so is $N(y = y')$ (resp., $N(y \neq y')$).[5]

Let $\mathcal{C}$ be a countably infinite set of constants not in Δ or φ such that there are countably many constants in $\mathcal{T}$ not in $\mathcal{C}$, Δ, or φ. (Our technical assumption ensures that such a sets $\mathcal{C}$ exists.)

Lemma 3.5: *If* $\mathbf{P}^+_\Lambda \not\vdash \Delta \hookrightarrow \varphi$*, then there exists a* $\mathcal{C}$*-good set* $\Delta^* \supseteq \Delta$ *such that* $\mathbf{P}^+_\Lambda \not\vdash \Delta^* \hookrightarrow \varphi^*$*, where* $\varphi^* = \varphi$ *if* $\varphi \in \mathcal{L}^{fo}$ *and* $\varphi^* = \varphi'[x_1/c_1, \ldots, x_k/c_k]$ *if* φ *has the form* $\forall x_1 \ldots \forall x_k \varphi'$*, where* φ' *is quantifier-free and* $c_1, \ldots, c_k$ *are constants in* $\mathcal{C}$*.*[6]

Once Lemma 3.5 is proved, the rest of the argument follows essentially the same lines as that of [Friedman, Halpern, and Koller 2000], so I just briefly outline the argument here. For each $c \in \mathcal{C}$, let $[c] = \{c' : c = c' \in \Delta^*\}$. We construct a model $M = (D, W, \pi, \mathcal{P})$ where $D = \{[c] : c \in \mathcal{C}\}$, W consists of all the $\mathcal{C}$-good sets Δ' such that $\Delta' \cap \mathcal{L}^0_C = \Delta^* \cap \mathcal{L}^0_C$ (so that worlds are $\mathcal{C}$-good sets), and π is such that, for all $\Delta' \in W$, we have $\pi(c, \Delta') = [c]$ and, for each atomic formula ψ, $(M, \Delta') \models \psi$ iff $\psi \in \Delta'$. This is consistent, since all sets Δ' in W agree on formulas of the form $N(c = c')$ and $N(c \neq c')$, and hence (since Δ' is $\mathcal{C}$-good) on formulas of the form

[5] This is an adaptation of the definition of $\mathcal{C}$-good given in [Friedman, Halpern, and Koller 2000], where the notion was applied to sets of formulas in $\mathcal{L}_C$ and a set $\mathcal{C}$ of constants.

[6] If the set Δ is restricted to being finite in formulas of the form $\Delta \hookrightarrow \psi$, then we replace the requirement $\mathbf{P}^+_\Lambda \not\vdash \Delta^* \hookrightarrow \varphi^*$ by $\mathbf{P}^+_\Lambda \not\vdash \Delta'' \hookrightarrow \varphi^*$ for all finite $\Delta'' \subseteq \Delta^*$. Analogous changes must be made throughout the proof.

$c = c$ and $c \neq c'$; thus, the formula $c = c'$ is in Δ' iff $\pi(c', \Delta') = \pi(c', \Delta') = [c]$. As shown in [Friedman, Halpern, and Koller 2000], it is possible to define a sequence $\mathcal{P}$ of probability measures such that $M \models \psi$ iff $\psi \in \Delta^*$ for all quantifier-free closed formulas $\psi \in \Delta^* \cap \mathcal{L}_C^0$ (note that if $\psi \in \Delta^* \cap \mathcal{L}_C^0$, then $\psi \in \Delta' \cap \mathcal{L}_C^0$ for all $\Delta' \in W$), and, if φ has the form $\forall x_1 \ldots \forall x_k \varphi'$, then $(M, \Delta^*) \models \neg(\varphi'[x_1/c_1, \ldots, x_1/c_k])$. This gives us the desired model.

So, to complete the proof of Theorem 3.4, it remains to prove Lemma 3.5. To do this, I construct Δ^* in stages. Starting with Δ, at each stage I add more and more formulas until I get a set Δ^* that is $\mathcal{C}$-good. Let $\mathcal{X}$ be the set of variables that appear in Δ and φ, and let $\mathcal{Y}$ be a countably infinite set of variables not in Δ or φ such that there are countably many variables not in $\mathcal{Y}$, Δ, or φ. (Again, our assumptions guarantee that such a set $\mathcal{Y}$ exists.) If $\varphi \in \mathcal{L}_C^0$, suppose that $\varphi = \forall x_1 \ldots \forall x_k \varphi'$. For convenience suppose that $\mathcal{Y}$ has the form $\{y_1, \ldots, y_k, z_1, z_2, z_3, \ldots\}$. (If $\varphi \in \mathcal{L}^{fo}$, then we can just take $\mathcal{Y} = \{z_1, z_2, z_3, \ldots\}$.) Let $\varphi^+ = \varphi$ if $\varphi \in \mathcal{L}^{fo}$ and $\varphi[x_1/y_1, \ldots, x_k/y_k]$ otherwise. Clearly

$$\mathbf{P}_\Lambda^+ \not\vdash \Delta \hookrightarrow \varphi^+. \tag{2}$$

This is immediate if $\varphi \in \mathcal{L}^{fo}$, since in that case $\varphi^+ = \varphi$. And if φ has the form $\forall x_1 \ldots \forall x_k \varphi'$ and $\mathbf{P}_\Lambda^+ \vdash \Delta \hookrightarrow \varphi'[x_1/y_1, \ldots, x_k/y_k]$, then it follows from Gen$^+$(b) that $\mathbf{P}_\Lambda^+ \vdash \Delta \hookrightarrow \varphi$, contradicting our assumption.

Let $\sigma_1, \sigma_2, \ldots$ be an enumeration of formulas in $\mathcal{L}^{fo} \cup \mathcal{L}_C^0$ over the vocabulary $\mathcal{T}$ whose only (free or bound) variables are in $\mathcal{Y}$. We define a sequence of set $\Delta_0, \Delta_1, \ldots$ inductively. Let $\Delta_0 = \Delta$. Suppose that $\Delta_1, \ldots, \Delta_m$ have been defined. Let $\exists x \psi$ be the $(m+1)$st existential first-order formula in the enumeration. Let $\Delta_{m+1} = \Delta_m \cup \{\exists x \psi \Rightarrow \psi[x/z_{m+1}]\}$ and let $\Delta^+ = \cup_m \Delta^m$. I claim that

$$\mathbf{P}_\Lambda^+ \not\vdash \Delta^+ \hookrightarrow \varphi^+. \tag{3}$$

If not, since $\mathbf{P}_\Lambda^+ \not\vdash \Delta \hookrightarrow \varphi$, there must exist some m such that $\mathbf{P}_\Lambda^+ \not\vdash \Delta^m \hookrightarrow \varphi$ and $\mathbf{P}_\Lambda^+ \vdash \Delta^{m+1} \hookrightarrow \varphi$. Since $\Delta^{m+1} = \Delta^m \cup \{\exists x \psi \Rightarrow \psi[x/z_{m+1}]\}$ and, by construction, z_{m+1} does not appear in Δ^m or φ, it follows from Gen$^+$(a) that $\mathbf{P}_\Lambda^+ \not\vdash \Delta^m \cup \{\exists x(\exists x \psi \Rightarrow \psi)\} \hookrightarrow \varphi$. It is easy to see that $\exists x(\exists x \psi \Rightarrow \psi)$ is a valid first-order formula. Thus, by IMP, $\mathbf{P}_\Lambda^+ \vdash \emptyset \hookrightarrow \exists x(\exists x \psi \Rightarrow \psi)$. It thus follows from TRANS that $\mathbf{P}_\Lambda^+ \vdash \Delta^m \hookrightarrow \varphi$, a contradiction.

I next construct a sequence $\Delta_0^+, \Delta_1^+, \ldots$ inductively, by taking $\Delta_0^+ = \Delta_0$, and taking $\Delta_{m+1}^+ = \Delta_m^+ \cup \{\sigma_{m+1}\}$ if $\mathbf{P}_\Lambda^+ \not\vdash \Delta_m^+ \cup \{\sigma_m\} \hookrightarrow \varphi$ and $\Delta_{m+1}^+ = \Delta_m^+$ otherwise. Let $\Delta^\dagger = \cup_m \Delta_m^+$. It is clear from the construction that $\mathbf{P}_\Lambda^+ \not\vdash \Delta_m^+ \hookrightarrow \varphi^+$ for all m; thus,

$$\mathbf{P}_\Lambda^+ \not\vdash \Delta^\dagger \hookrightarrow \varphi^+. \tag{4}$$

Several other properties of $\Delta^\dagger$ must be noted. First, the construction guarantees that if a first-order formula of the form $\exists x \psi$ is in $\Delta^\dagger$, then $\psi[x/y] \in \Delta^\dagger$ for some $y \in \mathcal{Y}$. It easily follows from F1$^+$ that if $\forall x \varphi \in \Delta^\dagger$ then $\varphi[x/y] \in \Delta^\dagger$ for all $y \in \mathcal{Y}$. Moreover, it follows from F5$^+$ that if $y \neq y' \in \Delta^\dagger$ for $y, y' \in \mathcal{Y}$ then $N(y \neq y') \in \Delta^\dagger$. Finally, it is easy to see that if $x = y \in \Delta^\dagger$, then $N(x = y) \in \Delta^\dagger$. (Proof: first observe that $\mathbf{P}_\Lambda^+ \vdash \emptyset \hookrightarrow (\mathit{false} \rightarrow \mathit{false})$ by REF and $\vdash_\Lambda x \neq x \Leftrightarrow \mathit{false}$, so $\mathbf{P}_\Lambda^+ \vdash \emptyset \hookrightarrow (x \neq x) \rightarrow \mathit{false}$, by LLE. Recall that, by definition, $(x \neq x) \rightarrow$

false $= N(x = x)$. By F4$^+$, $\mathbf{P}^+_\Lambda \vdash \{N(x = x), x = y\} \hookrightarrow N(x = y)$. By TRANS, $\mathbf{P}^+_\Lambda \vdash \{x = y\} \hookrightarrow N(x = y)$. It easily follows that if $x = y \in \Delta^\dagger$, then $N(x = y) \in \Delta^\dagger$.)

Let Δ^* be the result of replacing all occurrences of y_j in $\Delta^\dagger$ by c_j, for $j = 1, \ldots, k$ and replacing all occurrences of z_j in $\Delta^\dagger$ by d_j for $j = 1, 2, 3, \ldots$; similarly, let φ^* be the result of replacing all occurrences (if any) of y_j in φ^+ by c_j. It is easy to see that Δ^* is $\mathcal{C}$-good, and that $\mathbf{P}^+_\Lambda \not\vdash \Delta^* \hookrightarrow \varphi^*$ (if this is not the case, then it is immediate that (4) does not hold either). This completes the proof of Lemma 3.5 and, with it, the proof of Theorem 3.4. ∎

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 relate it to the more qualitative "asymptotic" 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 $\varphi \to \psi$, there are formulas of the form $\varphi \to^r \psi$, where r is a real number in $[0, 1]$. The semantics of such a formula is straightforward:

$(M, V) \models \varphi \to^r \psi$ 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 $\forall x_1 \ldots \forall x_n(\varphi \to^r \psi)$, where φ and ψ are first-order formulas.

Because $\mathcal{L}^0_{C,q}$ does not really consider limiting probability, it would be straightforward to give it semantics using a single probability measure Pr. That is, if a model M involves just a single distribution Pr rather than a sequence $(\Pr_1, \Pr_2, \ldots)$, then we could take $(M, \mathcal{V}) \models \varphi \to^r \psi$ if $\Pr([\![\psi]\!]_{M,V} \mid [\![\varphi]\!]_{M,V}) \geq 1 - r$. Indeed, that is essentially the semantics used in [Datta, Halpern, Mitchell, Roy, and Sen 2015], where the focus is on quantitative security. I continue to use a sequence of probability measures here, since my goal is to relate quantitative proofs to qualitative proofs.

That said, it is not hard to show that the same formulas are valid for the language $\mathcal{L}^0_{C,q}$ whether we consider a single probability measure or a sequence of probability measures in the semantics. This makes $\mathcal{L}^0_{C,q}$ closer to other approaches that have used probability and classical implication. A recent example is the work of Atserias and Balcázar [2015]. They consider formulas of the form $X \to Y$, where X and Y are sets of *attributes*, and can be identified with conjunctions of primitive propositions. Changing the notation of Atserias and Balcázar to be compatible with that of this paper, $\Pr \models_\gamma X \to Y$ if $\Pr(Y \mid X) \geq \gamma$. Given a set Δ of such formulas, $\Pr \models_\gamma \Delta \hookrightarrow (X \to Y)$ if $\Pr(Y' \mid X') \geq \gamma$ for all $X' \to Y' \in \Delta$ implies $\Pr(Y \mid X) \geq \gamma$. Atserias

and Balcázar provide algorithms for checking when such entailments hold, using linear programming. Note that in Atserias and Balcàzar's work, the same γ is used for the hypothesis and the conclusion in the entailment; it will be crucial for my main result that it is possible to use different parameters (the framework used here allows us to write the parameter explicitly in the formula).

I now now explain how qualitative reasoning in $\mathcal{L}^0_C$ and quantitative reasoning in $\mathcal{L}^0_{C,q}$ can be related. First note that for each of the axioms and rules in system $\mathbf{P}^+_\Lambda$, there is a corresponding sound axiom or rule in $\mathcal{L}^0_{C,q}$. Consider the following axioms:

LLE^q. $\{\varphi_1 \to^r \psi\} \hookrightarrow (\varphi_2 \to^r \psi)$ if $\vdash_\Lambda \varphi_1 \Leftrightarrow \varphi_2$.

RW^q. $\{\varphi \to^r \psi_1\} \hookrightarrow (\varphi \to^r \psi_2)$ if $\vdash_\Lambda \psi_1 \Rightarrow \psi_2$.

REF^q. $\emptyset \hookrightarrow (\varphi \to^0 \varphi)$.

AND^q. $\{\varphi \to^{r_1} \psi_1, \varphi \to^{r_2} \psi_2\} \hookrightarrow (\varphi \to^{r_3} (\psi_1 \wedge \psi_2))$, where $r_3 = \min(r_1 + r_2, 1)$.

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

CM^q. $\{\varphi_1 \to^{r_1} \varphi_2, \varphi_1 \to^{r_2} \psi\} \hookrightarrow ((\varphi_1 \wedge \varphi_2) \to^{r_3} \psi)$, where $r_3 = \min(r_1 + r_2, 1)$.

C6^q. $\{true \to^r false\} \hookrightarrow false$ for all $r \in [0, 1)$.

F5^q. $\{x \neq y\} \hookrightarrow N^0(x \neq y)$, where $N^0\varphi$ is an abbreviation for $\neg\varphi \to^0 false$.

Let $\mathbf{P}^{+,q}_\Lambda$ consist of the rules above, together with TRIV, F1$^+$, F4$^+$, IMP, Gen$^+$, and TRANS (all of which hold with no change in the quantitative setting), and

INC. $\{\varphi \to^{r_1} \psi\} \hookrightarrow (\varphi \to^{r_2} \psi)$ if $r_1 \leq r_2$ (increasing superscript).

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, Halpern, Mitchell, Roy, and Sen 2015] show, $\mathbf{P}^{+,q}_\Lambda$ (when combined with axioms for reasoning about actions and partial correctness) 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 $\Delta \subseteq \mathcal{L}^{fo} \cup \mathcal{L}^0_C$, say that $\Delta' \subseteq \mathcal{L}^{fo} \cup \mathcal{L}^0_{C,q}$ is a *quantitative instantiation* of Δ if there is a bijection f from Δ to Δ' such that, for every formula $\varphi \to \psi \in \Delta$, there is a real number $r \in [0, 1]$ such that $f(\varphi) = \varphi^r$, where $\varphi^r = \varphi$ if $\varphi \in \mathcal{L}^{fo}$, and $(\forall x_1 \ldots \forall x_k(\varphi' \to \psi))^r = \forall x_1 \ldots \forall x_k(\varphi' \to^r \psi)$. 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 $\Delta \hookrightarrow \varphi$ is derivable in $\mathbf{P}^+_\Lambda$ then, for all $r \in [0, 1]$, there exists a quantitative instantiation Δ' of Δ such that $\Delta' \hookrightarrow \varphi^r$ is

derivable in $\mathbf{P}^{+,q}_{\Lambda}$. Thus, if the system designer wants security at level r (that is, she wants to know that a 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 $\mathrm{AX}^{\Lambda}_{C}$), 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 $\mathbf{P}^{+}_{\Lambda} \vdash \Delta \hookrightarrow \varphi$, then for all $r \in [0,1]$, there exists a quantitative instantiation Δ' of a finite subset Δ'' of Δ such that $\mathbf{P}^{+,q}_{\Lambda} \vdash \Delta' \hookrightarrow \varphi^{r}$. Moreover, Δ' can be found in polynomial time, given the derivation of $\Delta \hookrightarrow \varphi$.*

Proof: Intuitively, Δ'' consists of the formulas in Δ needed for the proof of $\Delta \hookrightarrow \varphi$. The existence of Δ' follows by a straightforward induction on the length of the derivation. If it has length 1, then the proof must be an instance of an axiom. In this case, the argument proceeds by considering each axiom in turn. The arguments are all straightforward. I consider a few representative cases here. If the axiom TRIV was applied, then it must be the case that $\varphi \in \Delta$, so we can take $\Delta'' = \{\varphi\}$ and $\Delta' = \{\varphi^{r}\}$. If REF was applied, the φ has the form $\varphi' \rightarrow \varphi'$. In this case, take $\Delta'' = \emptyset$. By REF^{q}, $\emptyset \hookrightarrow (\varphi' \rightarrow^{0} \varphi')$. The conclusion now follows from INC. If AND was applied, then φ must have the form $\varphi' \rightarrow \psi_1 \wedge \psi_2$, where $\varphi' \rightarrow \psi_1, \varphi' \rightarrow \psi_2 \in \Delta$. Let $\Delta'' = \{\varphi' \rightarrow \psi_1, \varphi' \rightarrow \psi_2\}$. Choose $s_1, s_2 \in [0,1]$ such that $s_1 + s_2 = r$.[7] Let $\Delta' - \{\varphi' \rightarrow^{s_1} \psi_1, \varphi' \rightarrow^{s_2} \psi_2\}$. By AND^{q}, it easily follows that $\mathbf{P}^{+,q}_{\Lambda}\Delta' \hookrightarrow (\varphi \rightarrow^{r} (\psi_1 \wedge \psi_2))$. The argument for all the other axioms in $\mathbf{P}^{+}_{\Lambda}$ is similar and left to the reader.

Now suppose that the derivation of $\Delta \hookrightarrow \varphi$ has length $N > 1$. If the last line of the derivation is an axiom, the the argument above applies without change. Otherwise, it must be the result of applying Gen^{+}(a), Gen^{+}(b), or TRANS. I consider the latter two cases here; the case of Gen^{+}(a) is straightforward and left to the reader. If Gen^{+}(b) was applied, then φ has the form $\forall x \varphi'$ and there is a derivation $\Delta \hookrightarrow \varphi[x/z]$ of length less than N, where z does not appear in Δ. By the induction hypothesis, there is a subset Δ'' of Δ and instantiation of Δ' of Δ'' such that $\mathbf{P}^{+,q}_{\Lambda} \vdash \Delta' \hookrightarrow (\varphi')^{r}[x/z]$. Since $\varphi^{r} = (\forall x \varphi')^{r} = \forall x((\varphi')^{r})$, by Gen^{+}(a), it follows that $\mathbf{P}^{+,q}_{\Lambda} \vdash \Delta' \hookrightarrow \varphi^{r}$.

Finally, if the last step in the proof of $\Delta \hookrightarrow \varphi$ is an application of TRANS, then there exists $\Delta_1 \subseteq \Delta$ such that $\mathbf{P}^{+}_{\Lambda} \vdash \Delta \hookrightarrow \psi$ for all $\psi \in \Delta_1$ and $\mathbf{P}^{+}_{\Lambda} \vdash \Delta_1 \hookrightarrow \varphi$. By the induction hypothesis, there exists a finite subset Δ_2 of Δ_1 and a quantitative instantiation Δ'_2 of Δ_2 such that $\mathbf{P}^{+,q}_{\Lambda} \vdash \Delta'_2 \vdash \varphi^{r}$. By the induction hypothesis again, for each formula $\psi^{s} \in \Delta'_2$, there exists a finite subset Δ_{ψ} of Δ and a quantitative instantiation Δ'_{ψ} of Δ_{ψ} such that $\mathbf{P}^{+,q}_{\Lambda} \vdash \Delta'_{\psi} \hookrightarrow \psi^{s}$. Let $\Delta'' = \cup_{\psi \in \Delta_2} \Delta_{\psi}$ and let $\Delta' = \cup_{\psi \in \Delta_2} \Delta'_{\psi}$. By TRANS, we have $\mathbf{P}^{+,q}_{\Lambda}\Delta' \hookrightarrow \varphi^{r}$, as desired.

[7] 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.

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.

Given r, we might wonder how close to optimal the q we can find in Theorem 3.7 is. While I cannot give a definitive answer here, the following comments may give some insight. For each axiom individually, the quantitative instantiation is optimal, in the sense that it is not hard to find examples where the bounds are satisfied with equality. For example, in the AND^q rule, we cannot do better in general than taking $r_3 = \min(r_1 + r_2, 1)$. However, it could well be that when putting several steps in a proof together, we end up with a significant overestimate. Different proofs of the same conclusion might lead to different estimates. That also suggests that developing new proof rules might be useful, since they might result in better estimates.

4 Discussion

I have shown how the intuition behind the $\supset$ operator used by Datta et al. [2005] can be captured using the well-studied conditional implication operator $\rightarrow$, where $\varphi \rightarrow \psi$ is interpreted as "the probability of ψ conditional on φ converges to 1" (or converges to 1 super-polynomially). Using $\rightarrow$ with this interpretation has a significant advantage: it allows us to relate quantitative and qualitative reasoning. Specifically, for a rich fragment $\mathcal{L}_C^0$ of full first-order conditional logic, I have shown that if φ is provable from a collection Δ of formulas in $\mathcal{L}_C^0 \cup \mathcal{L}^{fo}$, then, for all r, there exists a quantitative instantiation Δ' of a finite subset of Δ, computable in time polynomial in the length of the proof of $\Delta \hookrightarrow \varphi$, such that φ^r is provable from Δ'. That mean that once we have a qualitative proof of a fact of interest from some assumptions, we can compute the strength of the assumptions needed to reach that conclusion.

I have suggested that this result is applicable to reasoning about quantitative security. This statement must be interpreted carefully. While $\mathcal{L}_C^0$ is a rich fragment of $\mathcal{L}_C$, it clearly does not suffice for stating and proving interesting facts about security programs. To do that, we need more operators than just those for reasoning about (conditional) probability. For example, the logic used in [Datta, Halpern, Mitchell, Roy, and Sen 2015] includes Hoare-like assertions of the form $\varphi[program]\psi$ and included axioms about the predictability of nonces and the unforgeability of signatures. Even if we had a completeness theorem for the fragment of the logic that does not involve

the $\rightarrow$ operator (which we do not), there may be interactions between the $\rightarrow$ operator and the other operators. So while the axioms of this paper are still sound, it is unlikely that they will be enough for completeness (even when they are added to complete axiomatizations for the other operators). The fact that (as shown by Theorem 3.1) that the language cannot distinguish super-polynomial convergence from non-super-polynomial convergence is further evidence that the logic cannot capture all that is needed for proving correctness of security protocols.

Nevertheless, that does not render the results of this paper irrelevant for security. Even in a richer logic, Theorem 3.7 still holds. More precisely, if there is a qualitative proof of a fact of the form $\Delta \hookrightarrow \varphi$ using only the axioms and rules of $\mathbf{P}^+_\Lambda$ for reasoning about $\rightarrow$ (or, more generally, using only axioms and rules that have quantitative analogues (as the axioms and rules of $\mathbf{P}^+_\Lambda$ do)), then, for all r, we can find a quantitative instantiation Δ' of a finite subset of Δ such that $\Delta' \hookrightarrow \varphi^r$ is provable. In [Datta, Halpern, Mitchell, Roy, and Sen 2015], nontrivial quantitative properties of a challenge-response protocol are proved using the axioms of $\mathbf{P}^{+,q}_\Lambda$ and axioms involving $\rightarrow^q$ that talk about properties of nonces and the likelihood that a signature can be forged. Interestingly, the superscript q used in the proof (and in the axioms that talk about the properties of nonces and unforgeability) is not a constant, but a function of the security parameter (and other parameters); this does not affect the arguments about $\rightarrow^q$ at all (all the axioms of this paper still hold), but again, makes it more applicable to security. It is easy to transform the quantitative proof given in [Datta, Halpern, Mitchell, Roy, and Sen 2015] to a qualitative proof of a conclusion of the form $\Delta \hookrightarrow \varphi$ where the only axioms used for reasoning about $\rightarrow$ are those in $\mathbf{P}^+_\Lambda$. We simply erase the superscript on $\rightarrow$, and put (qualitative versions of) the axioms for nonces and unforgeability into Δ. Now Theorem 3.7 can be used, for example, to deduce the strength of security parameter needed to get the conclusion to hold with a given strength. Put another way, we can think of the proof in [Datta, Halpern, Mitchell, Roy, and Sen 2015] as going in the "forward" direction: given assumptions about the security parameter and unforgeability, we prove that a conclusion about the security protocol holds with a certain probability. Theorem 3.7 lets us go in the "backward" direction: after proving a result qualitatively, we can deduce the strength that the assumptions need to hold with to be able to get the conclusion to hold with a given strength. It seems to me that both directions are of practical interest.

It is perhaps also worth noting that one of the features of the logic considered here is that it allows us to make statements about conditional probability; the formula $\varphi \rightarrow \psi$ is making a statement about the probability of φ conditional on ψ. Most of the security analyses in the literature have focused on unconditional probability; this amounts to considering formulas of the form $true \rightarrow \psi$. (This is the case for the analyses done by Datta et al. [2015], for example.) One reason for the focus on unconditional probability may be that the guarantees provided by cryptographic schemes are typically unconditional guarantees (or, at least, their formalizations avoid the use of unconditional probabilities). For example, when we say that a signature scheme has only a negligible probability of being broken, we take this probability to be unconditional (i.e., relative to the whole space). A signature scheme is considered secure even if someone broadcasts all the secret keys, as long as this is done with exponentially small probability. Of course, conditional on getting the secret keys, the

system is highly insecure. But we don't tend to worry about that in our proofs; it is a negligible event.

As shown by the analysis of Datta et al. [2015], the logic considered here is of interest even in the unconditional case (i.e., if we restrict to formulas of the form $true \rightarrow \psi$ or $true \rightarrow^r \varphi$). Moreover, although current analyses seem to focus on the unconditional case, it seems to me that a more refined security analysis might want to take into account some conditional probabilities: we might be interested, for example, in how secure a system would be if one of the agents in the system chose a somewhat weak key or inadvertently changed some security settings in a system. Note that once we allow nontrivial formulas φ in the antecedent of $\rightarrow$, it becomes critical that $\rightarrow$ is nonmonotonic. While a conclusion of the form $true \rightarrow \psi$ may be sound, a conclusion of the form $\varphi \rightarrow \psi$ may not be, if φ is an unlikely event (like an agent inadvertently changing some security setting).

Finally, it is worth considering in a little more detail the relationship between this work and that of Bana, Hasebe, and Okada [2013]. As I mentioned in the introduction, they also use a $\rightarrow$ operator. They give semantics to their $\rightarrow$ operator relative to a sequence $\mathcal{P}$ of probability measures, just as I do, but the the technical details of their semantics are quite different from the semantics I use.[8] Bana et al. work purely at the qualitative level; they have no "concrete" analogue $\rightarrow^r$ to their $\rightarrow$ operator. Thus, they make no attempt to relate their quatitative semantics to a more quantitative semantics. That said, they are very interested in relating qualitative work on what they call symbolic adversaries to more quantitative work on verification that works directly on models that involve probability. So, their goals are the same as mine. More experience is needed to determine which logic is better suited to reasoning about such programs. However it is done, as I hope the results of this paper have made clear, a logic with conditional statements can be an extremely useful addition to a security analyst's toolkit.

Acknowledgments: I thank Anupam Datta, John Mitchell, Riccardo Pucella, Arnab Roy, and Shayak Sen for many useful discussions on applying conditional logic to security protocols. The anonymous reviewers of the paper provided many helpful suggestions as well.

[8] In more detail, for Bana et al., each probability measure $\Pr_n$ is a measure on a possibly different domain W_n. They also consider sequences $S = (S_1, S_2, \ldots)$, where $S_n \subseteq W_n$, in giving the semantics of formulas. Very roughly speaking, $(\mathcal{P}, S) \models \varphi \rightarrow \psi$ if $(\mathcal{P}, S) \models \Box(p \Rightarrow q)$, where $\Box\varphi$ holds if φ holds for all *non-negligible* subsets S' of S, and $S' = (S'_1, S'_2, \ldots)$ is a non-negligible subset of S if $S'_n \subseteq S_n$ for all n and $\Pr_n(S'_n)$ is non-negligible as a function of n, that is, $\Pr_n(\Omega_n - S_n)$ does not grow super-polynomially.

There are clearly significant differences between the two semantics for $\rightarrow$. In the semantics for $\rightarrow$ that I have used, there is no analogue to the sequence S of events. Moreover, my focus is on events whose probability increases (super-polynomially) to 1, rather than non-negligible events. As an anonymous reviewer pointed out, we could make the two approaches somewhat closer by viewing the measures $\Pr_n$ used by Bana et al. as all being defined on a single space $\Omega = W_1 \times W_2 \times \cdots$ (i.e., the crossproduct of the domain of $\Pr_n$'s), where the measurable subsets of Ω are those of the form $A_1 \times A_2 \times \cdots$ and A_n is a measurable subset of W_n, defining $\Pr_n(A_1 \times A_2 \times \cdots) = \Pr_n(A_n)$, and identifying a sequence $S' = (S'_1, S'_2, \ldots)$ where $S'_n \subseteq W_n$ with the subset $S'_1 \times S_2 \times \cdots$ of Ω. But even with these identifications, the other differences pointed out above still remain.

References

Abadi, M. and P. Rogaway (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*, pp. 3–22. Springer-Verlag.

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

Atserias, A. and J. L. Balcázar (2015). Entailment among probabilistic relations. In *Proc. 30th IEEE Symposium on Logic in Computer Science*, pp. 621–632.

Bana, G., K. Hasebe, and M. Okada (2013). Computationally complete symbolic attacker and key exchange. In *Proc. 13th ACM SIGSAC Conference on Computer and Communications Security (CCS '13)*, pp. 1231–1246.

Barthe, G., B. Grégoire, S. Heraud, and S. Zanella-Béguelin (2011). Computer-aided security proofs for the working cryptographer. In P. Rogaway (Ed.), *Advances in Cryptology (CRYPTO 2011)*, Volume 6841 of *Lecture Notes in Computer Science*, pp. 71–90.

Barthe, G., B. Grégoire, and S. Zanella-Béguelin (2009). Formal certification of code-based cryptographic proofs. *SIGPLAN Notices 44*(1), 90–101.

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

Bellare, M., R. Canetti, and H. Krawczyk (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*.

Blanchet, B. (2006). A computationally sound mechanized prover for security protocols. In *IEEE Symposium on Security and Privacy*, pp. 140–154.

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

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

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

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

Datta, A., J. Y. Halpern, J. C. Mitchell, A. Roy, and S. Sen (2015). A symbolic logic with concrete bounds for cryptographic protocols. Available at http://arxiv.org/abs/1511.07536.

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

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

Friedman, N., J. Y. Halpern, and D. Koller (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., P. Morris, and J. Pearl (1993). A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions of Pattern Analysis and Machine Intelligence 15*(3), 220–232.

Goldszmidt, M. and J. Pearl (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)*, pp. 661–672.

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

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

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

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

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