A Logic for SDSI's Linked Local Name Spaces

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

Ron van der Meyden
School of Computer Science and Engineering,
University of New South Wales,
Sydney 2052,
Australia.
meyden@cse.unsw.edu.au
http://www.cse.unsw.edu.au/~meyden

January 19, 2000

Abstract

Abadi has introduced a logic to explicate the meaning of local names in SDSI, the Simple Distributed Security Infrastructure proposed by Rivest and Lampson. Abadi's logic does not correspond precisely to SDSI, however; it draws conclusions about local names that do not follow from SDSI's name resolution algorithm. Moreover, its semantics is somewhat unintuitive. This paper presents the Logic of Local Name Containment, which does not suffer from these deficiencies. It has a clear semantics and provides a tight characterization of SDSI name resolution. The semantics is shown to be closely related to that of logic programs, leading to an approach to the efficient implementation of queries concerning local names. A complete axiomatization of the logic is also provided.

1 Introduction

Rivest and Lampson [RL96] introduced SDSI—a Simple Distributed Security Infrastructure—to facilitate the construction of secure systems.[1] In SDSI, principals (agents) are identified with public keys. In addition to principals, SDSI allows other names, such as `poker-buddies`. Rather than having a global name space, these names are interpreted *locally*, by each principal. That is, each principal associates with each name a set of principals. Of course, the interpretation of a name such as `poker-buddies` may be different for each agent. However, a principal can "export" his bindings to other principals. Thus, Ron may receive a message from the principal he names `Joe` describing a set of principals Joe associates with `poker-buddies`. Ron may then refer to the principals Joe associates with `poker-buddies` by the expression `Joe's poker-buddies`.

Rivest and Lampson [RL96] give an operational account of local names; they provide a name-resolution algorithm that, given a principal `k` and a name `n`, computes the set of principals associated with `n` according to `k`. Abadi [Aba98] has provided a logic that, among other things, gives a more semantic account of local names. According to Abadi, its purpose "is to explain local names in a general, self-contained way, without requiring reference to particular implementations." Abadi shows that the SDSI name-resolution algorithm can be captured in terms of a collection of sound proof rules in his logic.

Abadi's focus is on axioms. He constructs a semantics, not with the goal of capturing the intended meaning of his constructs, but rather, with the goal of showing that certain formulas are not derivable from his axioms. (In particular, he shows that *false* is not derivable, showing that his axioms are consistent.) While adequate for Abadi's restricted goals, his semantics validates some formulas that we certainly would not expect to be valid. One consequence of this is that, while he is able to pinpoint some potential concerns with the logic, the resolution of these concerns is less satisfactory. For example, he observes that adding two seemingly reasonable axioms to his logic allows us to reach quite an unreasonable conclusion. However, it is not obvious from the semantic intuitions provided by Abadi which (if either) of the axioms is unreasonable, or why it is unreasonable. Moreover, while he proves that this particular unreasonable conclusion is not derivable in his framework, as we show, a closely related (and equally unreasonable) conclusion is in fact valid. This means we have no assurance that it or other similar formulas cannot be derived from Abadi's axioms.

We very much subscribe to Abadi's goal of using a logic to give a general account of naming. In this paper, we provide a logic whose syntax is very similar to Abadi's, but whose semantics is quite different and, we believe, captures better the meaning we intend the constructs to have. Nevertheless, all but one of Abadi's name space axioms are sound in our system.

We remark that, in a sense, our task is much easier than Abadi's, since we give the constructs in the logic a somewhat narrower reading than he does. Abadi tends to intertwine and occasionally identify issues of naming with issues of rights and delegation. (Such an identification is also implicitly made to some extent in designs such as PolicyMaker [BFL96].) We believe that it is important to treat these issues separately. Such a separation allows us to both

[1] SDSI now forms the basis for the Simple Public Key Infrastructure (SPKI) standardization work [Gro98]. SPKI simplifies some SDSI features (e.g., it eliminates groups) but adds many others. We focus in this paper on the core naming features of SDSI—there are some minor differences in the way that SPKI has chosen to handle these features, but we believe that our work is equally relevant to the the fragment of SPKI dealing with naming.

give a cleaner semantics for each of the relevant notions and to clarify a number of subtleties. This paper focuses on naming, which we carefully separate from the other issues; a companion paper [HvdMS99] considers authority and delegation.

We believe that our approach has a number of significant advantages:

- We can still simulate the SDSI's name resolution algorithm; Abadi's extra axiom is unnecessary. In fact, our logic captures SDSI's name resolution more accurately than Abadi's. Abadi's logic can draw conclusions that SDSI's name resolution cannot; our logic, in a precise sense, draws exactly the same conclusions as SDSI's name resolution algorithm.

- According to our semantic intuition, one of Abadi's proposed additional axioms is in fact quite unreasonable; it does not hold under our semantics, and it is quite clear why.

- We are able to provide a sound and complete axiomatization of our logic. Thus, unlike Abadi, we have a proof system that corresponds precisely to our semantics. This will allow us to prove stronger results than Abadi's about formulas that cannot be derived in our framework. Our completeness proof also yields a (provably optimal) NP-complete decision procedure for satisfiability of formulas in the logic.

- Our logic is closely related to Logic Programming. This allows us to translate queries about names to Logic Programming queries, and thus use all the well-developed Logic Programming technology to deal with such queries.

- Our approach opens the road to a number of generalizations, which allow us to deal with issues like permission, authority, and delegation [HvdMS99].

The rest of this paper is organized as follows. In Section 2, we review Abadi's logic and, in the process, describe SDSI's naming scheme. We also point out what we see as the problems with Abadi's approach. In Section 3, we give the syntax and semantics of our logic, and present a complete axiomatization. In Section 4 we show that our logic provides a tight characterization of SDSI name resolution. Section 5 deals with the connection between our account of SDSI name resolution and logic programming, and Section 6 concerns `Self`, an additional construct considered by Abadi. Section 7 concludes.

2 SDSI's Name Spaces and Abadi's Logic

In this section, we briefly review SDSI's naming scheme and Abadi's logic, and discuss our criticism of Abadi's logic. Like Abadi, we are basing our discussion on SDSI 1.1 [RL96].

2.1 SDSI's Name Spaces

SDSI has *local names* and a set of reserved names, which we refer to as *global names*. Both are associated with sets of principals, but the set of principals associated with a local name depends on the principal owning the local name space, while the set of principals associated with a global name does not. We denote the set of global names by G with generic element $\mathtt{g}$, the set of local names by N with generic element $\mathtt{n}$, and the set of keys (principals) by K with generic element $\mathtt{k}$. We assume that all these sets are pairwise disjoint and that K is nonempty. *Global identifiers* are either keys or global names.[2]

The elements of $K \cup G \cup N$ are said to be *simple names.* We form *principal expressions* from simple names inductively. Simple names are principal expressions, and if $\mathtt{p}$ and $\mathtt{q}$ are principal expressions, then so is ($\mathtt{p}$'s $\mathtt{q}$). Abadi's semantics (and ours) makes the latter operation associative, in that (($\mathtt{p}$'s $\mathtt{q}$)'s $\mathtt{r}$) and ($\mathtt{p}$'s ($\mathtt{q}$'s $\mathtt{r}$)) have the same meaning. In light of this, we can ignore parenthesization when writing such expressions. The expression $\mathtt{p}_1$'s $\ldots \mathtt{p}_{m-1}$'s $\mathtt{p}_m$ is written in SDSI as ($\mathtt{ref}$:$\mathtt{p}_1, \ldots, \mathtt{p}_m$).[3] We remark for future reference that SDSI has a special global name denoted "DNS!!", which represents the root of the DNS (Internet mail) hierarchy; this allows us to express an email address such as `bob@fudge.com` as DNS!!'s `com`'s `fudge`'s `bob`.

SDSI allows a principal to issue *certificates* of the form $\mathtt{n} \longmapsto \mathtt{p}$, signed with its key. If $\mathtt{k}$ issues such a certificate, it has the effect of binding local name $\mathtt{n}$ in $\mathtt{k}$'s name space to the principals denoted by the principal expression $\mathtt{p}$.[4] Notice that only principals issue certificates, and that these certificates bind a local name (not a global name) to some set of principals. In general, a local name may be bound to a unique principal, no principal, or many principals. SDSI allows a principal $\mathtt{k}$ to issue certificates $\mathtt{n} \longmapsto \mathtt{p}_1$ and $\mathtt{n} \longmapsto \mathtt{p}_2$. This has the effect of binding $\mathtt{n}$ to (at least) the principals denoted by $\mathtt{p}_1$ and $\mathtt{p}_2$.

SDSI provides a name-resolution algorithm for computing the set of principals bound to a name. The core of the algorithm consists of a nondeterministic procedure REF2. For ease of exposition, we take REF2 to have four arguments: a principal $\mathtt{k}$, a function c that associates with each principal $\mathtt{k}'$ a set of bindings (intuitively, ones that correspond to certificates signed by $\mathtt{k}'$), a function β which associates with each global name $\mathtt{g}$ a set of principals (intuitively, the ones bound to $\mathtt{g}$), and a principal expression $\mathtt{p}$. REF2($\mathtt{k}$,β,c,$\mathtt{p}$) returns the principal(s) bound to $\mathtt{p}$ in $\mathtt{k}$'s name space, given the bindings β and the certificates c. REF2 is nondeterministic; the set of possible outputs of REF2 is taken to be the set of principals bound to $\mathtt{p}$ in $\mathtt{k}$'s name space. REF2 is described in Figure 1.[5]

2.2 Abadi's Logic: Syntax, Semantics, and Axiomatization

The formulas in Abadi's logic are formed by starting with a set of primitive propositions and formulas of the form $\mathtt{p} \longmapsto \mathtt{p}'$, where $\mathtt{p}$ and $\mathtt{p}'$ are principal expressions. More complicated for-

[2]Note that Abadi uses G for global identifier; thus, his G corresponds to our $G \cup K$.

[3]SDSI allows m to be 0, taking (`ref:`) to be the current principal. In Section 6, we follow Abadi by considering an expression `Self` that represents (`ref:`).

[4]SDSI also allows other forms of binding that we do not consider here. Our notation is also a simplification of that used by SDSI.

[5]Our version of REF2 is similar, although not identical, to Abadi's. Like Abadi's, it is simpler than that in [RL96], in that we do not deal with a number of issues, such as quoting or encrypted objects, dealt with by SDSI. Our presentation of REF2 differs from Abadi's mainly in its treatment of global names. Abadi assumes that REF2 takes only two arguments, $\mathtt{o}$ and $\mathtt{p}$, where $\mathtt{o}$ is either a global identifier (i.e., an element of $G \cup K$) or `current principal`, denoted `cp`. Although he does not write c explicitly as an argument, he does assume that there is a set he denotes assumptions($\mathtt{o}$) that includes bindings corresponding to signed certificates. In addition, it includes bindings for `cp`. We do not have a distinguished current principal; rather, if the current principal is $\mathtt{k}$, then for uniformity we assume that all of the current principal's bindings are also described by the bindings in $c(\mathtt{k})$. More significantly, if $\mathtt{g}$ is a global name, then Abadi's REF2($\mathtt{o}$,$\mathtt{g}$) would return $\mathtt{g}$, while ours would return some principal $\mathtt{k}$ to which $\mathtt{g}$ is bound in β. Our approach seems more consistent with the SDSI presentation of REF2, but this difference is minor, and all of Abadi's results hold for our presentation of REF2.

```
REF2(k,β,c,p)
   if p ∈ K then return(p)
   else if p ∈ G
        then if β(p) = ∅ then fail
             else return(k′) for some k′ ∈ β(p)
   else if p is a local name n in N
        then if c(k) = ∅ then fail
             else for some n ⟼ q ∈ c(k) return(REF2(k,β,c,q))
   else if p is of the form q's r
        then return(REF2(REF2(k,β,c,q),β,c,r))
```

Figure 1: Procedure REF2

mulas are formed by closing off under conjunction, negation, and formulas of the form $\mathtt{p}$ *says* ϕ, where ϕ is a formula.

Abadi views $\mathtt{p} \longmapsto \mathtt{p}'$ as meaning that $\mathtt{p}$ is "bound to" $\mathtt{p}'$. He considers two possible interpretations of "bound to". The first is equality; however, he rejects this as being inappropriate. (In particular, it does not satisfy some of his axioms.) The second is that $\mathtt{p} \longmapsto \mathtt{p}'$ means $\mathtt{p}'$ "speaks-for" $\mathtt{p}$, in the sense discussed in [ABLP93, LABW92]. Roughly speaking, this says that any message certified by $\mathtt{p}'$ should be viewed as also having been certified by $\mathtt{p}$. While the "speaking-for" interpretation is the one favored by Abadi, he does not commit to it. Note that under Abadi's "speaking-for" interpretation, it makes sense to write $\mathtt{p} \longmapsto \mathtt{p}'$ for arbitrary principal expressions $\mathtt{p}$ and $\mathtt{p}'$. However, SDSI allows only local (simple) names to be bound to principal expressions. We shall make a similar restriction in our logic (and, indeed, under our semantic interpretation of binding, it would not make sense to allow an arbitrary principal expression to be bound to another one.)

The "speaks-for" interpretation intertwines issues of delegation with those of naming. As we suggested in the introduction, we believe these issues should be separated. We shall give $\longmapsto$ a different interpretation that we believe is simpler and more in the spirit of binding. We believe that the "speaks-for" relation of [ABLP93, LABW92] should have quite different semantics than that of binding names to principals. (We hope to return to this issue in future work.)

Abadi interprets $\mathtt{p}$ *says* ϕ as "the principal denoted by p makes a statement that implies ϕ". In the case where $\mathtt{p}$ is a key (i.e., principal) $\mathtt{k}$, this could mean that $\mathtt{k}$ signs a statement saying ϕ. Under our more restrictive interpretation, this is exactly how we interpret our analogue to *says*.

In any case, note that Abadi translates SDSI's local name $\mathtt{n}$ being bound to $\mathtt{p}$ as $\mathtt{n} \longmapsto \mathtt{p}$ and captures $\mathtt{k}$ signing a certificate saying $\mathtt{n}$ is bound to $\mathtt{p}$ by the formula $\mathtt{k}$ *says* $\mathtt{n} \longmapsto \mathtt{p}$. For future reference, it is worth noting that, in order to capture the binding of names to principals, no use is made of primitive propositions.

Abadi interprets formulas in his logic with respect to a tuple $(\mathcal{W}, \alpha, \rho, \mu)$. The function α maps global identifiers ($G \cup K$) to subsets of $\mathcal{W}$. The function ρ maps $N \times \mathcal{W}$ to subsets of $\mathcal{W}$.

Finally, μ associates with each world (principal) $\mathtt{k}$ and primitive proposition p a truth value $\mu(p, \mathtt{k})$.

Abadi does not provide any intuition for his semantics, but suggests that $\mathcal{W}$ should be thought of as a set of possible worlds, as in modal logic. However, he also suggests [private communication, 1999] that his semantics was motivated by the work of Grove and Halpern [GH93], in which the corresponding set contains *pairs* consisting of a world and an agent. Some of Abadi's definitions make more intuitive sense if we think of $\mathcal{W}$ as a set of agents, while others make more sense if we think of $\mathcal{W}$ as a set of worlds. We elaborate on this point below.

Given $\mathtt{k} \in \mathcal{W}$ and $\mathtt{p} \in P$, Abadi defines $[\![\mathtt{p}]\!]_{\mathtt{k}}$ inductively, as follows:

- $[\![\mathtt{g}]\!]_{\mathtt{k}} = \alpha(\mathtt{g})$, for $\mathtt{g} \in G \cup K$
- $[\![\mathtt{n}]\!]_{\mathtt{k}} = \rho(\mathtt{n}, \mathtt{k})$ for $\mathtt{n} \in N$
- $[\![\mathtt{p}_1\text{'s } \mathtt{p}_2]\!]_{\mathtt{k}} = \cup\{[\![\mathtt{p}_2]\!]_{\mathtt{k}'} : \mathtt{k}' \in [\![\mathtt{p}_1]\!]_{\mathtt{k}}\}$

Here we have used a notation corresponding to the interpretation of the "worlds" in $\mathcal{W}$ as agents. Using this interpretation we may think of $[\![\mathtt{p}]\!]_{\mathtt{k}}$ as the set of principals bound to principal expression $\mathtt{p}$ according to $\mathtt{k}$. The clause for $[\![\mathtt{p}_1\text{'s } \mathtt{p}_2]\!]_{\mathtt{k}}$ then says that if $\mathtt{k}'$ is one of the principals referred to by $\mathtt{k}$ as $\mathtt{p}_1$, then $\mathtt{k}$ uses $\mathtt{p}_1$'s $\mathtt{p}_2$ to refer to any principal referred to by $\mathtt{k}'$ as $\mathtt{p}_2$.

Abadi also defines what it means for a formula ϕ to be true at world $\mathtt{k} \in \mathcal{W}$, written $\mathtt{k} \models \phi$, inductively, by

- $\mathtt{k} \models p$ iff $\mu(p, \mathtt{k}) = \textbf{true}$, if p is a primitive proposition
- $\mathtt{k} \models \phi \wedge \psi$ iff $\mathtt{k} \models \phi$ and $\mathtt{k} \models \psi$
- $\mathtt{k} \models \neg\phi$ iff $\mathtt{k} \not\models \phi$
- $\mathtt{k} \models \mathtt{p} \longmapsto \mathtt{p}'$ iff $[\![\mathtt{p}]\!]_{\mathtt{k}} \subseteq [\![\mathtt{p}']\!]_{\mathtt{k}}$
- $\mathtt{k} \models \mathtt{p}$ *says* ϕ iff $\mathtt{k}' \models \phi$ for all $\mathtt{k}' \in [\![\mathtt{p}]\!]_{\mathtt{k}}$.

These clauses defining $\models$ are quite intuitive if one interprets $\mathcal{W}$ to be a set of worlds and considers $[\![\mathtt{p}]\!]_{\mathtt{k}}$ to be the set of worlds consistent with what principal $\mathtt{p}$ has said at world $\mathtt{k}$. In particular, under this interpretation, the clause for *says* can be read as stating that $\mathtt{p}$ *says* ϕ if ϕ holds in all worlds consistent with what $\mathtt{p}$ has said. The clause for $\longmapsto$ also has quite a plausible reading under the "speaks-for" interpretation of this construct: it states that $\mathtt{p}'$ speaks for $\mathtt{p}$ if all worlds consistent with what $\mathtt{p}$ has said are consistent with what $\mathtt{p}'$ has said, i.e., $\mathtt{p}$ is constrained to speak consistently with what $\mathtt{p}'$ has said. However, it seems rather difficult to extend this intuitive reading to encompass the inductive definition of $[\![\mathtt{p}]\!]_{\mathtt{k}}$. In particular, it is far from clear to us what intuitive understanding to assign to the clause for $[\![\mathtt{p}_1\text{'s } \mathtt{p}_2]\!]_{\mathtt{k}}$ on this reading.

On the other hand, note that if we interpret the worlds as agents, then we can think of $\mathtt{k} \models \phi$ as saying that ϕ is true when local names are interpreted according to agent $\mathtt{k}$. But this reading of the clauses, when combined with the intuitive reading of $[\![\mathtt{p}]\!]_{\mathtt{k}}$ as the set of principals that $\mathtt{k}$ refers to using $\mathtt{p}$, also has its difficulties. Intuitively, when $\mathtt{n}$ is bound to $\mathtt{p}$ in principal

Reflexivity:	$\mathtt{p} \longmapsto \mathtt{p}$
Transitivity:	$(\mathtt{p} \longmapsto \mathtt{q}) \Rightarrow ((\mathtt{q} \longmapsto \mathtt{r}) \Rightarrow (\mathtt{p} \longmapsto \mathtt{r}))$
Left Monotonicity:	$(\mathtt{p} \longmapsto \mathtt{q}) \Rightarrow ((\mathtt{p}\text{'s } \mathtt{r}) \longmapsto (\mathtt{q}\text{'s } \mathtt{r}))$
Globality:	$(\mathtt{p}\text{'s } \mathtt{g}) \longmapsto \mathtt{g}$ if $\mathtt{g}$ is a global identifier
Associativity:	$((\mathtt{p}\text{'s } \mathtt{q})\text{'s } \mathtt{r}) \longmapsto (\mathtt{p}\text{'s } (\mathtt{q}\text{'s } \mathtt{r}))$
	$(\mathtt{p}\text{'s } (\mathtt{q}\text{'s } \mathtt{r})) \longmapsto ((\mathtt{p}\text{'s } \mathtt{q})\text{'s } \mathtt{r})$
Linking:	$(\mathtt{p}\ \mathit{says}\ (\mathtt{n} \longmapsto \mathtt{r}) \Rightarrow ((\mathtt{p}\text{'s } \mathtt{n}) \longmapsto (\mathtt{p}\text{'s } \mathtt{r}))$ if $\mathtt{n}$ is a local name
Speaking-for:	$(\mathtt{p} \longmapsto \mathtt{q}) \Rightarrow ((\mathtt{q}\ \mathit{says}\ \phi) \Rightarrow \mathtt{p}\ \mathit{says}\ \phi)$

Figure 2: Abadi's axioms for linked local name spaces

k's local name space, the principals that k refers to using p should be a subset of the principals that k refers to using n. Abadi interprets n being bound to p as $\mathtt{n} \longmapsto \mathtt{p}$; this holds with respect to principal k when $[\![\mathtt{p}]\!]_{\mathtt{k}}$ is a *superset* of $[\![\mathtt{n}]\!]_{\mathtt{k}}$. This is precisely the opposite of what we would expect. Thus, neither the interpretation of $\mathcal{W}$ as a set of worlds nor the interpretation of $\mathcal{W}$ as a set of agents gives a fully satisfactory justification for Abadi's semantics. As we shall see, in our semantics, the interpretation of a principal expression p according to an agent will be a set of agents, but we use the reverse of Abadi's containment to represent binding.

Abadi provides an axiom system for his logic, which has three components:

1. The standard axioms and rules of propositional logic.

2. The standard axiom and rule for modal logic for the *says* operator:

$$(\mathtt{p}\ \mathit{says}\ (\phi \Rightarrow \psi)) \Rightarrow ((\mathtt{p}\ \mathit{says}\ \phi) \Rightarrow (\mathtt{p}\ \mathit{says}\ \psi))$$

$$\frac{\phi}{\mathtt{p}\ \mathit{says}\ \phi}$$

3. New axioms dealing with linked local name spaces, shown in Figure 2.

He shows that this axiomatization is sound, but conjectures it is not complete.

2.3 Name Resolution in Abadi's Logic

Abadi proves a number of interesting results relating his logic to SDSI. First, he shows that in a precise sense his logic can simulate REF2. He provides a collection of name-resolution rules *NR* and proves the following results:[6]

Proposition 2.1: *Given a collection of c of bindings corresponding to signed certificates and a set β of bindings of global names to keys, let E be the conjunction of the formulas* $\mathtt{k}\ \mathit{says}\ \mathtt{n} \longmapsto \mathtt{q}$

[6]The results stated here are a variant of those stated in Abadi's paper, since our version of REF2 differs slightly from his. Nevertheless, the proofs of the results are essentially identical.

for each certificate $\texttt{n} \longmapsto \texttt{q} \in c(\texttt{k})$ *and the formulas* $\texttt{g} \longmapsto \texttt{k}$ *for each* $\texttt{k} \in \beta(\texttt{g})$. *Then* $E \Rightarrow ((\texttt{k}\text{'s } \texttt{p}) \longmapsto \texttt{k}_1)$ *is provable with the name resolution rules NR if and only if* $\text{REF2}(\texttt{k}, \beta, c, \texttt{p})$ *yields* $\texttt{k}_1$.

Proposition 2.2: *The name resolution rules are sound with respect to the logic. That is, given* E *as in Proposition 2.1 and any principal expression* $\texttt{p}$, *if* $E \Rightarrow (\texttt{p} \longmapsto \texttt{k})$ *is provable using NR then* $E \Rightarrow (\texttt{p} \longmapsto \texttt{k})$ *is also provable in the logic.*

These results show that any bindings of names to principals that can be deduced using REF2 can also be deduced using Abadi's logic. However, Abadi shows that his logic is actually more powerful than REF2, by giving two examples of conclusions that can be deduced from his logic but not using REF2:

Example 2.3: Using the Globality, Associativity, and Transitivity axioms, if $\texttt{k}$ and $\texttt{k}'$ are keys, we immediately get $\texttt{k}$'s ($\texttt{Lampson}$'s $\texttt{k}'$) $\longmapsto \texttt{k}'$. This result does not follow from the REF2 algorithm. That is, $\text{REF2}(\texttt{k}, \beta, c, \texttt{Lampson}\text{'s } \texttt{k}')$ does not necessarily yield $\texttt{k}'$ for arbitrary c and β (in particular, it will not do so if $\texttt{Lampson}$ is not bound to anything in c). ∎

Example 2.4: Suppose c consists of the four certificates that correspond to the following formulas: $\texttt{k}$ *says* $(\texttt{Lampson} \longmapsto \texttt{k}_1)$, $\texttt{k}$ *says* $(\texttt{Lampson} \longmapsto \texttt{k}_2)$, $\texttt{k}_1$ *says* $(\texttt{Ron} \longmapsto \texttt{Rivest})$, and $\texttt{k}_2$ *says* $(\texttt{Rivest} \longmapsto \texttt{k}_3)$ (where $\texttt{k}$, $\texttt{k}_1$, $\texttt{k}_2$, and $\texttt{k}_3$ are keys). Using the Speaking-for axiom, it is not hard to show that we can conclude that $\texttt{k}$'s ($\texttt{Lampson's Ron}$) $\longmapsto \texttt{k}_3$. It is easy to show that REF2 cannot reach this conclusion; that is, $\text{REF2}(\texttt{k}, \beta, c, \texttt{Lampson's Ron})$ does not yield $\texttt{k}_3$ for any β.[7] ∎

In reference to Example 2.3, Abadi [Aba98] says that "it is not clear whether [these conclusions] are harmful, and they might in fact be useful". In general, he views it as a feature of his logic that it allows reasoning about names without knowing their bindings [private communication, 1999]. While we agree that, in general, reasoning about names without knowing their bindings is a powerful feature, we believe it is important to make clear exactly which conclusions are desirable and which are not. This is what a good semantics can provide. Under our semantics, neither of these two conclusions are valid. In fact, our logic draws precisely the same conclusions as REF2. Of course, the conclusions of Examples 2.3 and 2.4 are valid under Abadi's semantics but, as we observed earlier, Abadi's semantics is not really meant to be used as a guide to which conclusions are acceptable (and, indeed, as we shall see, it validates a number of conclusions that do not seem so acceptable).

Abadi also considers the effect of extending his axiom system. In particular, he considers adding the following two axioms:

- the converse of Globality: $\texttt{g} \longmapsto (\texttt{p}\text{'s } \texttt{g})$

[7]SPKI certificates and SDSI certificates have a slightly different syntactic form. A SPKI certificate issued by $\texttt{k}$ to bind $\texttt{n}$ to $\texttt{p}$ could be expressed in the logic as $\texttt{k}$ *says* ($\texttt{k}$'s $\texttt{n} \longmapsto \texttt{p}$). Abadi has remarked [private communication 1999], that if we rewrite the example using assertions in this form, the corresponding conclusion of this example would not follow in his logic. We have followed the SDSI format for certificates in this paper, but note that after some minor changes to the definitions, all the results in Sections 3–5 would still apply to SPKI certificates.

- a generalization of Linking: $(\mathtt{p}\ \mathit{says}\ (\mathtt{p}_1 \longmapsto \mathtt{p}_2)) \Rightarrow (\mathtt{p}\text{'s } \mathtt{p}_1 \longmapsto \mathtt{p}\text{'s } \mathtt{p}_2)$, for an arbitrary principal $\mathtt{p}_1$ (instead of a local name).

The generalization of Linking is in fact sound under Abadi's semantics. The converse of Globality is not, but only because we may have $[\![\mathtt{p}]\!]_{\mathtt{k}} = \emptyset$. Note that $[\![\mathtt{p}]\!]_{\mathtt{k}} = \emptyset$ iff $\mathtt{k} \models \mathtt{p}\ \mathit{says}\ \mathit{false}$; thus, the following variant of the converse of Globability is sound under Abadi's semantics: $\neg(\mathtt{p}\ \mathit{says}\ \mathit{false}) \Rightarrow (\mathtt{g} \longmapsto (\mathtt{p}\text{'s } \mathtt{g}))$.

This is quite relevant to our purposes because Abadi shows that if we added the two axioms above to his system, then from $\mathtt{k}\ \mathit{says}\ (\mathtt{DNS!!} \longmapsto \mathtt{k})$, we can conclude $\mathtt{DNS!!} \longmapsto \mathtt{k}$. Thus, just from k saying that DNS!! is bound to k, it follows that DNS!! is indeed bound to k. This is particularly disconcerting under Abadi's "speaks-for" interpretation, where $\mathtt{DNS!!} \longmapsto \mathtt{k}$ becomes "k speaks for DNS!!". We certainly do not want an arbitrary principal to speak for the name server!

Abadi proves a result showing that such conclusions are not derivable from hypotheses of a certain type in his logic (which does not have these two axioms).

Proposition 2.5: [Aba98] *Let* $\mathtt{k}$ *and* $\mathtt{k}'$ *be distinct global names; let* ϕ *be a formula of the form* $(\mathtt{k}'\ \mathit{says}\ (\mathtt{n}_1 \longmapsto \mathtt{p}_1)) \wedge \ldots \wedge (\mathtt{k}'\ \mathit{says}\ (\mathtt{n}_k \longmapsto \mathtt{p}_k))$*, where* $\mathtt{n}_1, \ldots, \mathtt{n}_k$ *are local names and* $\mathtt{p}_1, \ldots, \mathtt{p}_k$ *are principal expressions; let* ψ *be a formula of the form* $(\mathtt{k}\ \mathit{says}\ \psi_1) \wedge \ldots \wedge (\mathtt{k}\ \mathit{says}\ \psi_m)$*, where* $\psi_1, \ldots, \psi_m$ *are arbitrary formulas. Then* $\phi \wedge \psi \Rightarrow (\mathtt{k}' \longmapsto \mathtt{k})$ *is not valid.*[8]

While Proposition 2.5 provides some assurance that undesirable formulas are not derivable in the logic, it does not provide much. Indeed, if we allow the ψ to include the formula $\neg(\mathtt{k}'\ \mathit{says}\ \mathit{false})$, then the result no longer holds. In fact, it follows from our earlier discussion that the formula

$$(\mathtt{k}\ \mathit{says}\ (\mathtt{DNS!!} \longmapsto \mathtt{k})) \wedge \neg(\mathtt{k}\ \mathit{says}\ \mathit{false}) \Rightarrow (\mathtt{DNS!!} \longmapsto \mathtt{k})$$

is valid. Moreover, it does not seem so unreasonable to allow conjuncts such as $\neg(\mathtt{k}\ \mathit{says}\ \mathit{false})$ as part of ψ. We certainly want to be able to use the logic to be able to say that if a principal's statements are not blatantly inconsistent, then certain conclusions follow.

3 The Logic of Local Name Containment

In this section we propose the Logic of Local Name Containment (henceforth LLNC) as an alternative to Abadi's logic. LLNC interprets local names as sets of principals and interprets SDSI certificates as stating containment relationships between these sets. We define the syntax in Section 3.1. In Section 3.2 we describe two distinct semantics for the logic. Section 3.3 presents a complete axiomatization.

3.1 Syntax

LLNC has syntactic elements that are closely related to the syntactic elements of Abadi's logic. However, our notation differs slightly from Abadi's to help emphasize some of the differences in intuition.

[8] Abadi's result actually says "$\phi \wedge \psi \Rightarrow (\mathtt{k}' \longmapsto \mathtt{k})$ is not derivable"; since his axiomatization is sound, but not necessarily complete, the claim that it is not valid is stronger, and that is what Abadi's proof shows.

Again, we start with keys K, global names G, and local names N, and form principal expressions from them. The formulas of our language are formed as follows:

- If $\mathtt{p}$ and $\mathtt{q}$ are principal expressions then $\mathtt{p} \longmapsto \mathtt{q}$ is a formula.
- If $\mathtt{k} \in K$ and ϕ is a formula then $\mathtt{k}$ *cert* ϕ is a formula.[9]
- If ϕ_1 and ϕ_2 are formulas, then so are $\neg\phi_1$ and $\phi_1 \wedge \phi_2$. As usual, $\phi_1 \vee \phi_2$ is an abbreviation for $\neg(\neg\phi_1 \wedge \neg\phi_2)$ and $\phi_1 \Rightarrow \phi_2$ is an abbreviation for $\neg\phi_1 \vee \phi_2$.

We write $\mathcal{L}$ for the set of all formulas. (For simplicity, we omit primitive propositions, although we could easily add them. They play no role in Abadi's account of SDSI names, nor will they in ours.)

We read the expression $\mathtt{p} \longmapsto \mathtt{q}$ as "$\mathtt{p}$ contains $\mathtt{q}$"; we intend for it to capture the fact that all the keys bound to $\mathtt{q}$ are also bound to $\mathtt{p}$. However, our intuitions about the meaning of $\mathtt{p} \longmapsto \mathtt{q}$ are quite different from Abadi's. In particular, we do not wish to interpret $\mathtt{p} \longmapsto \mathtt{q}$ as "$\mathtt{q}$ speaks for $\mathtt{p}$." We consider the "speaks for" relation as being about rights and delegation, which requires a more sophisticated semantics than we wish to consider here. (See [HvdMS99] for a logic for reasoning about rights and delegation.) The expression $\mathtt{p} \longmapsto \mathtt{q}$ should be understood as simply asserting a containment relationship between the denotations of principal expressions $\mathtt{p}$ and $\mathtt{q}$; this is exactly what our semantics will enforce.

We read the expression $\mathtt{k}$ *cert* ϕ as "$\mathtt{k}$ has certified that ϕ." This corresponds roughly to Abadi's $\mathtt{k}$ *says* ϕ. There are two significant differences, however. For one thing, we do not allow arbitrary principal expressions on the left-hand side; only keys may certify a formula ϕ. For another, our interpretation of *cert* is more restrictive than Abadi's *says*, in that *cert* is treated quite syntactically; it refers to an actual certificate issued by a principal, while *says* considers logical consequences of such certificates. As a consequence, whereas *says* satisfies standard properties of modal operators (e.g., closure under logical consequence), *cert* does not.

3.2 Semantics

Our semantics is designed to model the SDSI principle that principals bind names in their local name space to values by issuing certificates. The interpretation of a local name depends on the principal and the certificates that have been issued. As the principal may rely on others for its interpretation of local names, the certificates issued by other principals also play a role. The interpretation of global names and keys will be independent of both the principal and the certificates that have been issued.

A *world* is a pair $w = (\beta, c)$, where $\beta : G \to \mathcal{P}(K)$ and $c : K \to \mathcal{P}(\mathcal{L})$ (where $\mathcal{P}(X)$ denotes the set of subsets of X) and $\cup_{\mathtt{k}\in K} c(\mathtt{k})$ is finite. Intuitively, the function β interprets global (or fixed) names as sets of keys. The intended interpretation of the function c is that it associates

[9]For our account of SDSI naming, it would suffice to restrict this clause to formulas of the form $\mathtt{k}$ *cert* $\mathtt{n} \longmapsto \mathtt{p}$ where $\mathtt{n} \in N$ and $\mathtt{p} \in P$: our semantics will treat more general certificates as irrelevant to the meaning of principal expressions. We allow the more general form for purposes of discussion and because we envisage generalizations of the logic in which other types of certificates will be required.

with every key $\mathtt{k}$ the set of formulas $c(\mathtt{k})$ that have been certified using this key. That is, if $\phi \in c(\mathtt{k})$ then, intuitively, a certificate asserting ϕ has been signed using $\mathtt{k}$.[10]

Formulas of the logic will be interpreted in a world with respect to a key. Intuitively, this key indicates the principal from whose perspective we interpret principal expressions.

To interpret local names, we introduce an additional semantic construct. A *local name assignment* will be a function $l : K \times N \to \mathcal{P}(K)$ associating each key and local name with a set of keys. Intuitively, $l(\mathtt{k}, \mathtt{n})$ is the set of keys represented by principal $\mathtt{k}$'s local name $\mathtt{n}$. We write *LNA* for the set of all local name assignments.

Given a world $w = (\beta, c)$, a local name assignment l, and a key $\mathtt{k}$, we may assign to each principal expression $\mathtt{p}$ an interpretation $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}}$, a set of keys. The definition is much like that of Abadi's $[\![p]\!]_{\mathtt{k}}$:

- $[\![\mathtt{k}']\!]_{w,l,\mathtt{k}} = \{\mathtt{k}'\}$, if $\mathtt{k}' \in K$ is a key,
- $[\![\mathtt{g}]\!]_{w,l,\mathtt{k}} = \beta(\mathtt{g})$, if $\mathtt{g} \in G$ is a global name,
- $[\![\mathtt{n}]\!]_{w,l,\mathtt{k}} = l(\mathtt{k}, \mathtt{n})$, if $\mathtt{n} \in N$ is a local name,
- $[\![\mathtt{p}\text{'s } \mathtt{q}]\!]_{w,l,\mathtt{k}} = \bigcup\{[\![\mathtt{q}]\!]_{w,l,\mathtt{k}'} \mid \mathtt{k}' \in [\![\mathtt{p}]\!]_{w,l,\mathtt{k}}\}$, for principal expressions $\mathtt{p}, \mathtt{q} \in P$.

Our intuitions for $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}}$ are essentially the same as for the "agent-based" reading of Abadi's logic, discussed above. That is, $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}}$ is the set of keys associated with the expression $\mathtt{p}$ in $\mathtt{k}$'s local name space, when local names are interpreted according to l. With respect to principal $\mathtt{k}$, the expression $\mathtt{p}$'s $\mathtt{q}$ denotes the set of principals that principals referred to by $\mathtt{k}$ as $\mathtt{p}$ refer to as $\mathtt{q}$.

We now define what it means for a formula ϕ to be true at a world $w = (\beta, c)$ with respect to a local name assignment l and key $\mathtt{k}$, written $w, l, \mathtt{k} \models \phi$, by induction on the structure of ϕ.[11]

- $w, l, \mathtt{k} \models \mathtt{p} \longmapsto \mathtt{q}$ if $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}} \supseteq [\![\mathtt{q}]\!]_{w,l,\mathtt{k}}$
- $w, l, \mathtt{k} \models \mathtt{k}'$ *cert* ϕ if $\phi \in c(\mathtt{k}')$
- $w, l, \mathtt{k} \models \neg\phi_1$ if not $w, l, \mathtt{k} \models \phi_1$
- $w, l, \mathtt{k} \models \phi_1 \wedge \phi_2$ if $w, l, \mathtt{k} \models \phi_1$ and $w, l, \mathtt{k} \models \phi_2$.

Note that the semantics of *cert* reinforces its syntactic nature. To determine if $\mathtt{k}'$ *cert* ϕ is true at $(w, l, \mathtt{k})$, we check whether a certificate has been issued in world w by $\mathtt{k}'$ certifying ϕ. Moreover, as we shall see, while we allow any formula to be certified by $\mathtt{k}$, the only formulas whose certification has a nontrivial semantic impact are those of the form $\mathtt{n} \longmapsto \mathtt{p}$, where $\mathtt{n}$ is a local name. We return to this issue below.

[10] We make the simplifying assumption that certificates do not have expiration dates. It is not difficult to extend the logic to take into account certificate expiration; see [HvdM99]. The assumption that $\cup_{\mathtt{k} \in K} c(\mathtt{k})$ is finite is meant to enforce the intuition that only finitely many certificates are issued. None of our later results depend on this assumption, but it seems reasonable given the intended application of the logic.

[11] Note that our semantics is thus in the spirit of that of Grove and Halpern [GH93], in that the truth of a formula depends on both an agent and some features of the world (captured by w and l).

We do not consider all pairs w, l as being appropriate on the left-hand side of $\models$. If $w = (\beta, c)$, we expect the local name assignment l to respect the certificates that have been issued in c. That is, if $c(\mathtt{k})$ includes the binding $\mathtt{n} \longmapsto \mathtt{p}$, we would expect that $l(\mathtt{k}, \mathtt{n})$ would include all the keys bound to $\mathtt{p}$ in $\mathtt{k}$'s name space. The question is whether there can be other keys bound to $\mathtt{n}$ in $\mathtt{k}$'s name space beyond those forced by the certificates. How we answer this question depends on our intuitions for c. For example, we could view c as the set of certificates received by one of the principals. This would be particularly appropriate if we wanted to reason about the knowledge and belief of the agents, an extension we plan to explore in future work. With this viewpoint, we could view l as consisting of all the bindings, including ones that the principal does not know about. Thus, l would at least have all the bindings forced by c, but perhaps others as well. Alternatively, we could view c as consisting of all the certificates that have been issued. In this case, we would want l to be in some sense *minimal*, and have no bindings beyond those forced by the certificates in c. We now present two different semantics, which reflect each of these two intuitions. We then show that, as far as validity is concerned, the semantics are equivalent; that is, they have the same proof theory.

A local name assignment l is *consistent* with a world $w = (\beta, c)$ if, for all keys $\mathtt{k}$, local names $\mathtt{n}$, and principal expressions $\mathtt{p}$, if the formula $\mathtt{n} \longmapsto \mathtt{p}$ is in $c(\mathtt{k})$, then $w, l, \mathtt{k} \models \mathtt{n} \longmapsto \mathtt{p}$. Intuitively, assignments that are not consistent with a world provide an inappropriate basis for the interpretation of local names, since the certificates issued by principals are not necessarily reflected in their local bindings. We obtain our first semantics, called the *open* semantics, by restricting to consistent local name assignments. We write $w, l, \mathtt{k} \models_o \phi$ if $w, l, \mathtt{k} \models \phi$ and l is consistent with w. The formula ϕ is *o-satisfiable* if there exists a triple $w, l, \mathtt{k}$ such that $w, l, \mathtt{k} \models_o \phi$ and ϕ is *o-valid*, denoted $\models_o \phi$, if there does not exist a triple $w, l, \mathtt{k}$ such that $w, l, \mathtt{k} \models_o \neg\phi$.

Although our syntax allows $\mathtt{k}$ to certify arbitrary formulas, it is easy to see that, according to the semantics just introduced (as well as the one we are about to introduce), only the certification of formulas of the form $\mathtt{n} \longmapsto \mathtt{p}$ has any impact on consistency; all other formulas certified by $\mathtt{k}$ are ignored. There is a good reason for this restriction. We are implicitly assuming that when $\mathtt{k}'$ certifies $\mathtt{n} \longmapsto \mathtt{p}$, that very act causes all the keys bound to $\mathtt{p}$ to also be bound to $\mathtt{n}$ in $\mathtt{k}$'s name space. Thus, if $\mathtt{n} \longmapsto \mathtt{p} \in c(\mathtt{k})$, then we want $\mathtt{n} \longmapsto \mathtt{p}$ to be true in $(w, l, \mathtt{k})$. But if $\mathtt{k}$ certifies a formula like $\mathtt{k}_1$'s $\mathtt{n} \longmapsto \mathtt{k}_3$ where $\mathtt{k}_1 \neq \mathtt{k}$, then we cannot conclude that this formula is true in $(w, l, \mathtt{k})$ unless we are prepared to make additional assumptions about $\mathtt{k}$'s truthfulness. We feel that if such assumptions are to be made, then they should be modeled explicitly in the logic, not hidden in the semantics.

It does seem reasonable to extend the notion of l being consistent with w to require that if $\mathtt{k}$ certifies a formula ψ which is a Boolean combination of formulas of the form $\mathtt{n} \longmapsto \mathtt{p}$ then $(w, l, \mathtt{k}) \models \psi$. However, once we allow more general Boolean combinations (in particular, once we allow disjunctions), there will be problems making sense out of the intuition of our next semantics, that there are "no bindings beyond those forced by the certificates in c". We consider this issue next.

According to the open semantics, it is possible for a local name $\mathtt{n}$ of principal $\mathtt{k}_1$ to be bound to a key $\mathtt{k}_2$ even when no certificate concerning $\mathtt{n}$ has been issued. Arguably, this is not in accordance with the intentions of SDSI. To better capture these intentions, we define a second semantics, that restricts the name bindings to those forced by the certificates issued.

To do so, we first establish that the open semantics satisfies a kind of "minimal model" result. Define the ordering $\leq$ on the space LNA of local name assignments by $l_1 \leq l_2$ if $l_1(\mathtt{k},\mathtt{n}) \subseteq l_2(\mathtt{k},\mathtt{n})$ for all $\mathtt{k} \in K$ and $\mathtt{n} \in N$. It is readily seen that LNA is given the structure of a complete lattice [Bir67] by this relation. Say that a local name assignment l is *minimal* in a set of local name assignments L if $l \in L$ and $l \leq l'$ for all $l' \in L$.

Theorem 3.1: *Given a world w, there exists a unique local name assignment l_w minimal in the set of all local name assignments consistent with w. Moreover, if $\mathtt{p}$ is a principal expression and $\mathtt{k}_1$ and $\mathtt{k}_2$ are keys, then $w, l_w, \mathtt{k}_1 \models_{\mathrm{o}} \mathtt{p} \longmapsto \mathtt{k}_2$ iff, for all local name assignments l consistent with w, we have $w, l, \mathtt{k}_1 \models_{\mathrm{o}} \mathtt{p} \longmapsto \mathtt{k}_2$.*

The proof of this result (which, like that of all the technical results in this paper, is deferred to the appendix) uses standard techniques from the theory of fixed points.

We now define our second semantics, called the *closed semantics.* It attempts to capture the intuition that the only bindings in l should be those required by the certificates in c, using the minimal assignment promised by Theorem 3.1. We write $w, \mathtt{k} \models_{\mathrm{c}} \phi$ if $w, l_w, \mathtt{k} \models \phi$. We say that ϕ is *c-satisfiable* if there exists a world w and key $\mathtt{k}$ such that $w, \mathtt{k} \models_{\mathrm{c}} \phi$ and that ϕ is *c-valid*, denoted $\models_{\mathrm{c}} \phi$, if $w, \mathtt{k} \models_{\mathrm{c}} \phi$ for all worlds w and principals $\mathtt{k}$. Note that by Theorem 3.1, the assignment l_w is consistent with w, so c-satisfiability implies o-satisfiability. Thus, if $\models_{\mathrm{o}} \phi$ then $\models_{\mathrm{c}} \phi$. As we shall soon see (Theorem 3.5), somewhat surprisingly, the converse holds as well.

3.3 A Complete Axiomatization

We start this section by presenting a sound and complete axiomatization for $LLNC$ with respect to the open semantics. We then prove that the open and closed semantics are characterized by the same valid formulas, so that the axiomatization is also sound and complete with respect to the closed semantics.

The axiomatization depends in part on whether the set K of keys is finite or infinite. Figure 3 describes the axiom system $\mathrm{AX}_{\mathit{inf}}$ for the case where K is infinite.

It is interesting to compare the axioms in $\mathrm{AX}_{\mathit{inf}}$ to Abadi's axioms. Although we interpret $\longmapsto$ as superset and he interprets it as subset, Reflexivity, Transitivity, Left-Monotonicity, and Associativity, hold in both cases, for essentially the same reasons. The switch from subset to superset means that the Converse of Globality holds in our case. Globality does not hold in general because the denotation of $\mathtt{p}$'s $\mathtt{g}$ may be empty if the denotation of $\mathtt{p}$ is empty (as we observed, this is also why the Converse of Globality does not hold in general for Abadi). In fact, for our logic, $\mathtt{p}$'s $\mathtt{g} \longmapsto \mathtt{g}$ holds whenever the interpretation of $\mathtt{p}$ is nonempty. We use $\mathtt{p}$'s $\mathtt{k} \longmapsto \mathtt{k}$ as a canonical way of denoting that the interpretation of $\mathtt{p}$ is nonempty. This explains the form of the Globality axiom. Since the interpretation of a key is always nonempty, we also get Key Globality.

Key Linking is our analogue of Abadi's Linking axiom. Of course, we use *cert* whereas Abadi uses *says*; in addition, only keys can certify formulas for us. While this axiom shows that there are some similarities between *cert* and *says*, there are some significant differences. We have no analogue of Abadi's Speaking-for axiom and, unlike *says*, *cert* does not satisfy the standard axiom and rule of modal logic: $(\mathtt{k}\ \mathit{cert}\ (\phi \Rightarrow \psi)) \wedge (\mathtt{k}\ \mathit{cert}\ \phi)$ does not imply $\mathtt{k}\ \mathit{cert}\ \psi$

Propositional Logic:	All instances of propositional tautologies
Reflexivity:	$\mathtt{p} \longmapsto \mathtt{p}$
Transitivity:	$(\mathtt{p} \longmapsto \mathtt{q}) \Rightarrow ((\mathtt{q} \longmapsto \mathtt{r}) \Rightarrow (\mathtt{p} \longmapsto \mathtt{r}))$
Left Monotonicity:	$(\mathtt{p} \longmapsto \mathtt{q}) \Rightarrow ((\mathtt{p}\text{'s } \mathtt{r}) \longmapsto (\mathtt{q}\text{'s } \mathtt{r}))$
Associativity:	$((\mathtt{p}\text{'s } \mathtt{q})\text{'s } \mathtt{r}) \longmapsto (\mathtt{p}\text{'s } (\mathtt{q}\text{'s } \mathtt{r}))$
	$(\mathtt{p}\text{'s } (\mathtt{q}\text{'s } \mathtt{r})) \longmapsto ((\mathtt{p}\text{'s } \mathtt{q})\text{'s } \mathtt{r})$
Key Globality:	$(\mathtt{k}\text{'s } \mathtt{g}) \longmapsto \mathtt{g}$ if $\mathtt{k} \in K$ and $\mathtt{g} \in G \cup K$
Globality:	$(\mathtt{p}\text{'s } \mathtt{k} \longmapsto \mathtt{k}) \Rightarrow (\mathtt{p}\text{'s } \mathtt{g} \longmapsto \mathtt{g})$ if $\mathtt{k} \in K,\ \mathtt{g} \in G \cup K$
Converse of Globality:	$\mathtt{g} \longmapsto (\mathtt{p}\text{'s } \mathtt{g})$ if $\mathtt{g} \in K \cup G$
Key Linking:	$(\mathtt{k}\ \mathit{cert}\ (\mathtt{n} \longmapsto \mathtt{r})) \Rightarrow ((\mathtt{k}\text{'s } \mathtt{n}) \longmapsto (\mathtt{k}\text{'s } \mathtt{r}))$
	if $\mathtt{n}$ is a local name
Nonemptiness:	(a) $\mathtt{p} \longmapsto \mathtt{k}_1 \Rightarrow \mathtt{p}\text{'s } \mathtt{k} \longmapsto \mathtt{k}$
	(b) $\neg(\mathtt{p} \longmapsto \mathtt{q}) \Rightarrow \mathtt{q}\text{'s } \mathtt{k} \longmapsto \mathtt{k}$
	(c) $\mathtt{p}\text{'s } \mathtt{q} \longmapsto \mathtt{k}_1 \Rightarrow \mathtt{p}\text{'s } \mathtt{k} \longmapsto \mathtt{k}$
	(d) $(\mathtt{p}\text{'s } \mathtt{k} \longmapsto \mathtt{k} \wedge \mathtt{k}' \longmapsto \mathtt{p}) \Rightarrow (\mathtt{p} \longmapsto \mathtt{k}')$
Key Distinctness:	$\neg(\mathtt{k}_1 \longmapsto \mathtt{k}_2)$ if $\mathtt{k}_1$ and $\mathtt{k}_2$ are distinct keys
Modus Ponens:	From ϕ and $\phi \Rightarrow \psi$ infer ψ.

Figure 3: The axiom system $\mathrm{AX}_{\mathit{inf}}$

and $\mathtt{k}$ *cert* ϕ is not valid even if ϕ is valid. Interestingly, Abadi does not use these properties of Speaking-for in proving that his name resolution rules *NR*, used to capture REF2, are sound. As a result, (with very minor changes) we can show that the name resolution rules are also sound for LLNC, and hence we can prove analogues of Propositions 2.1 and 2.2. However, we can actually prove a much stronger result: whereas Abadi's logic is able to draw conclusions about bindings that do not follow from REF2, LLNC captures REF2 exactly (see Theorem 4.1).

$\mathrm{AX}_{\mathit{inf}}$ has two axioms that do not appear in Abadi's axiomatization: Key Distinctness and Nonemptiness. Key Distinctness just captures the fact that we interpret keys as themselves. The first three parts of Nonemptiness capture various ways that an expression can be seen to be nonempty. For example, part (a) says that if $\mathtt{p}$ is bound to (i.e., is a superset of) a key, then its interpretation must be nonempty and part (b) says that if $\mathtt{p}$ is not a superset of $\mathtt{q}$, then $\mathtt{q}$ must be nonempty. Part (d) of Nonemptiness says that if $\mathtt{p}$ is nonempty and $\mathtt{k}'$ is bound to $\mathtt{p}$, then $\mathtt{p}$ is bound to $\mathtt{k}'$, i.e., $\mathtt{p}$ and $\mathtt{k}'$ have exactly the same interpretation.

If K is finite we need to add two further axioms to $\mathrm{AX}_{\mathit{inf}}$. Let $\mathrm{AX}_{\mathit{fin}}$ consist of all the axioms and rule in $\mathrm{AX}_{\mathit{inf}}$ together with:

Witnesses:	$\neg(\mathtt{p} \longmapsto \mathtt{q}) \Rightarrow \vee_{\mathtt{k} \in K}(\neg(\mathtt{p} \longmapsto \mathtt{k}) \wedge (\mathtt{q} \longmapsto \mathtt{k}))$
	$(\mathtt{p}\text{'s } \mathtt{q}) \longmapsto \mathtt{k}_1 \Rightarrow \vee_{\mathtt{k} \in K}((\mathtt{p} \longmapsto \mathtt{k}) \wedge (\mathtt{k}\text{'s } \mathtt{q} \longmapsto \mathtt{k}_1))$
Current Principal:	$\vee_{\mathtt{k} \in K}(\mathtt{n}_\mathtt{k} \longmapsto \mathtt{l}_\mathtt{k} \Leftrightarrow \mathtt{k}\text{'s } \mathtt{n}_\mathtt{k} \longmapsto \mathtt{l}_\mathtt{k})$
	where $\mathtt{n}_\mathtt{k} \in N$ and $\mathtt{l}_\mathtt{k} \in K$ for each $\mathtt{k} \in K$.

The two axioms that make up Witnesses essentially capture our interpretation of $\longmapsto$ as containment. They tell us that facts about containment of principal expressions can be reduced to facts about keys. For example, the first one says that if $\mathtt{p}$ does not contain $\mathtt{q}$, then there is a key bound to $\mathtt{q}$ that is not bound to $\mathtt{p}$. Current Principal captures the fact that some key in K must be the current principal; if $\mathtt{k}$ is the current principal, then for all local names $\mathtt{n}$ and keys $\mathtt{k}'$, $\mathtt{n} \longmapsto \mathtt{k}' \Leftrightarrow \mathtt{k}$'s $\mathtt{n} \longmapsto \mathtt{k}'$ holds. (This is actually true not just for local names, but for all principal expressions; it suffices to state the axiom just for local names.)

While the properties captured by these two axioms continue to hold even if K is infinite, they can no longer be expressed in the logic, since we cannot take a disjunction over all the elements in K. Interestingly, we can drop Nonemptiness and Globality as axioms in $\mathrm{AX}_{\mathit{fin}}$. These properties already follow from the other properties in the presence of Witnesses.

As the following result shows, these axiom systems completely characterize validity in the logic with respect to the open semantics.

Theorem 3.2: *AX_{inf} (resp., AX_{fin}) is a sound and complete axiomatization of LLNC with respect to the open semantics if K is infinite (resp., K is finite).*

In the course of proving Theorem 3.2, we also prove a "finite model" result, which we cull out here. Let $|\phi|$, the *length* of ϕ, be the total number of symbols appearing in ϕ. This result holds both when K is finite and when K is infinite.

Proposition 3.3: *Let K_ϕ be the keys that appear in ϕ and let $C_\phi(\mathtt{k})$ consist of all bindings $\mathtt{n} \longmapsto \mathtt{p}$ such that $\mathtt{k}$ cert $\mathtt{n} \longmapsto \mathtt{p}$ is a subformula of ϕ. If ϕ is satisfiable with respect to the open semantics, then for all sets K' of keys such that $K_\phi \subseteq K'$ and $|K'| \geq \min(|K|, 2 \cdot |\phi|^2)$, there is a world $w = (\beta, c)$, local name assignment l, and principal $\mathtt{k} \in K'$ such that $w, l, \mathtt{k} \models_{\mathrm{o}} \phi$ and (a) $l(\mathtt{k}', \mathtt{n}) \subseteq K'$ for all $\mathtt{k}' \in K$ and $\mathtt{n} \in N$, (b) $l(\mathtt{k}', \mathtt{n}) = \emptyset$ if $\mathtt{k}' \notin K'$, (c) $\beta(\mathtt{g}) \subseteq K'$ for all $\mathtt{g} \in G$, (d) $\beta(\mathtt{g}) = \emptyset$ if $\mathtt{g}$ does not occur in ϕ, and (e) $c(\mathtt{k}) \subseteq C_\phi(\mathtt{k})$ for all keys $\mathtt{k}$.*

Corollary 3.4: *The problem of deciding if a formula $\phi \in LLNC$ is satisfiable with respect to the open semantics is NP-complete (whether K is finite or infinite).*

Proof: The lower bound is immediate from the fact that we can trivially embed satisfiability for propositional logic into satisfiability for LLNC. For the upper bound, given ϕ, choose K' such that $|K'| = \min(|K|, 2 \cdot |\phi|^2)$ and $K' \supseteq K_\phi$. Then guess $w, l, \mathtt{k}$ as in Proposition 3.3 and check whether $w, l, \mathtt{k} \models_{\mathrm{o}} \phi$. Proposition 3.3 says that the guess is only polynomial in $|\phi|$; it is clear that checking whether $w, l, \mathtt{k} \models_{\mathrm{o}} \phi$ can also be done in time polynomial in ϕ. Note that for $|\phi| \leq |K|$ (which is likely to include all cases of practical interest, given that K will typically be a very large set), the polynomial does not depend on $|K|$. ∎

As we suggested earlier, the closed semantics and the open semantics are characterized by exactly the same axioms.

Theorem 3.5: *The same formulas are c-valid and o-valid; i.e., for all formulas ϕ, we have $\models_{\mathrm{o}} \phi$ iff $\models_{\mathrm{c}} \phi$.*

We remark that this result is sensitive to the language under consideration. It may no longer hold if we move to a more expressive language.

Corollary 3.6: *AX_{inf} (resp., AX_{fin}) is a sound and complete axiomatization of LLNC with respect to the closed semantics when K is infinite (rep., finite).*

Corollary 3.7: *The problem of deciding if a formula $\phi \in LLNC$ is satisfiable with respect to the closed semantics is NP-complete (whether K is finite or infinite).*

Let us now return to the contentious axioms discussed by Abadi. Converse of Globality is valid in $LLNC$, as we observed earlier. The generalization of Linking considered by Abadi, restricted to be syntactically well formed, amounts to

$$(\mathtt{k}\ \mathit{cert}\ (\mathtt{p}_1 \longmapsto \mathtt{p}_2)) \Rightarrow (\mathtt{k}\text{'s}\ \mathtt{p}_1 \longmapsto \mathtt{k}\text{'s}\ \mathtt{p}_2).$$

In general, this is not valid, since our semantics ignores certificates stating $\mathtt{p}_1 \longmapsto \mathtt{p}_2$ when $\mathtt{p}_1$ is not a local name. Thus, we avoid the "unreasonable" conclusions that can be drawn from these axioms. In particular, it does not follow in our logic that $(\mathtt{k}\ \mathit{cert}\ (\mathtt{DNS!!} \longmapsto \mathtt{k})) \Rightarrow \mathtt{DNS!!} \longmapsto \mathtt{k}$. However, the reason it does not follow in LLNC is quite different from the reason it does not follow in Abadi's logic: since DNS!! is a global name, a certificate such as $\mathtt{k}\ \mathit{cert}\ (\mathtt{DNS!!} \longmapsto \mathtt{k})$ has no impact on the interpretation of global names. This captures the intuition that k should not be trusted when making assertions about bindings not under its control. If we were willing to trust k on everything, then concluding that k is bound to DNS!! after k certifies that it is would not seem so unreasonable.

The following formula is also not valid in LLNC:

$$(\neg(\mathtt{k}\ \mathit{cert}\ \mathit{false}) \wedge (\mathtt{k}\ \mathit{cert}\ (\mathtt{DNS!!} \longmapsto \mathtt{k}))) \Rightarrow \mathtt{DNS!!} \longmapsto \mathtt{k}.$$

(This formula corresponds to the one that we noted earlier is valid in Abadi's logic.) Failure to issue a certificate stating *false* has no more impact on global names than does any other behavior of k. Nor would a precondition asserting that the interpretation of k is non-empty validate the formula, since this is true in every world. We can in fact prove the following generalization of Abadi's Proposition 2.5, which provides a stronger statement of the safety of our logic than Abadi's result.

Proposition 3.8: *Let Γ be any c-satisfiable boolean combination of formulas of the form $\mathtt{k}\ \mathit{cert}\ \phi$, and let Δ be any boolean combination of formulas of the form $\mathtt{p} \longmapsto \mathtt{q}$ where neither $\mathtt{p}$ nor $\mathtt{q}$ contains a local name. Then $\models_c \Gamma \Rightarrow \Delta$ iff $\models_c \Delta$.*

Informally, Proposition 3.8 says that facts about global names are completely independent of facts about certificates; issuing certificates can have no impact on the global name assignment. As we observed earlier, the analogous result does *not* hold for Abadi's logic.

4 Name Resolution in LLNC

In this section, we show that LLNC captures REF2 exactly. Indeed, we show that it does so for several distinct semantic interpretations. Define the order $\geq$ on worlds by $(\beta', c') \geq (\beta, c)$ if

1. $\beta'(\mathtt{g}) \supseteq \beta(\mathtt{g})$ for all global names $\mathtt{g}$, and

2. $c'(\mathtt{k}) \supseteq c(\mathtt{k})$ for all keys $\mathtt{k}$.

That is, $w' \geq w$ when w' contains more certificates than w and the bindings to global names in w are a subset of those in w'. If E is a set of formulas and ϕ is a formula, we write $E \models_{\mathrm{o}} \phi$ if for all worlds w, local name assigments l consistent with w and all keys $\mathtt{k}$, if $w, l, \mathtt{k} \models_{\mathrm{o}} \psi$ for all ψ in E then $w, l, \mathtt{k} \models_{\mathrm{o}} \phi$. Similarly, $E \models_{\mathrm{o}} \phi$ if for all worlds w and all keys $\mathtt{k}$, if $w, \mathtt{k} \models_{\mathrm{c}} \psi$ for all ψ in E then $w, \mathtt{k} \models_{\mathrm{c}} \phi$.

Theorem 4.1: *Suppose $\mathtt{k}_1, \mathtt{k}_2$ are principals, $w = (\beta, c)$ is a world, and $\mathtt{p}$ is a principal expression. Let E_w be the set of all formulas $\mathtt{g} \longmapsto \mathtt{k}$ for all global names $\mathtt{g}$ and keys $\mathtt{k} \in \beta(\mathtt{g})$ and the formulas $\mathtt{k}$ cert ϕ for all keys $\mathtt{k}$ and formulas $\phi \in c(\mathtt{k})$. The following are equivalent:*

1. $\mathtt{k}_1 \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{p})$,

2. $w, \mathtt{k}_2 \models_{\mathrm{c}} \mathtt{p} \longmapsto \mathtt{k}_1$,

3. $w', \mathtt{k}_2 \models_{\mathrm{c}} \mathtt{p} \longmapsto \mathtt{k}_1$ *for all worlds* $w' \geq w$,

4. $E_w \models_c \mathtt{k}_2\text{'s } \mathtt{p} \longmapsto \mathtt{k}_1$,

5. $E_w \models_o \mathtt{k}_2\text{'s } \mathtt{p} \longmapsto \mathtt{k}_1$.

This theorem gives a number of perspectives on name resolution in LLNC. The equivalence between (1) and (2) in this theorem tells us that REF2 is sound and complete with respect to key binding, according to the semantics of LLNC. That is, $\mathrm{REF2}(\mathtt{k}, \beta, c, \mathtt{p})$ yields $\mathtt{k}'$ iff $\mathtt{p} \longmapsto \mathtt{k}'$ is forced to be true by the bindings of global names in β and the certificates in c. Thus, viewed as a specification of the meaning of SDSI names, the closed semantics and REF2 are equivalent.

Informally, we have viewed REF2 as a procedure that is run by an omniscient agent with complete information about the interpretation of global names and the certificates that have been issued. It is also possible to understand REF2 as performing a computation based on the limited information available to a particular principal. Suppose that the world w expresses the limited information this principal has about the binding of global names and the certificates that have been issued. Suppose that w' describes the actual bindings of global names and the certificates that have been issued. Assuming that all of the principal's information is correct, then $w \leq w'$. Thus, the set of $w' \geq w$ is the set of all worlds w' that are consistent with the information available to the principal. (We could formalize this using the Kripke semantics for the logic of knowledge in a distributed system [HM90].) The equivalence between (2) and (3) essentially shows that it doesn't matter whether we view the principal as having total or partial information.

The implication from (1) to (4) in Theorem 4.1 is analogous to Abadi's soundness result, Proposition 2.2. Of course, the converse implication gives us completeness, which, as Abadi himself observed, does not hold for Abadi's logic (since it validates conclusions that do not follow from REF2). Interestingly, although, as we have seen, there are significant differences between *LLNC* and Abadi's logic, an examination of Abadi's soundness proof reveals that it does not use the Speaking-for rule, the unrestricted form of Globality, or the standard axiom and rule for the modal operator *says*, which are the main points of difference with our logic. This observation says that the proof of the implication from (1) to (4) is essentially the same for $LLNC$ and for Abadi's logic.

It is instructive to understand why the formulas considered in Examples 2.3 and 2.4, which give conclusions in Abadi's logic beyond those derivable by REF2, are not valid in LLNC. It is easy to see why the formula $\texttt{k}$'s ($\texttt{Lampson}$'s $\texttt{k}'$) $\longmapsto \texttt{k}'$ from Example 2.3 (which, by Associativity and Transitivity, is equivalent to ($\texttt{k}$'s $\texttt{Lampson}$)'s $\texttt{k}' \longmapsto \texttt{k}'$) is not valid in LLNC. This is simply because the antecedent of (our version of) Globality does not always hold. Now consider the formula in Example 2.4. The proof that this is valid in Abadi's logic uses the Speaking-for axiom, which does not hold for us (if we replace *says* by *cert*). To see that it is not valid in LLNC, consider a world $w = (\beta, c)$ containing only the certificates forced by the formulas (i.e., $c(\texttt{k}) = \{\texttt{Lampson} \longmapsto \texttt{k}_1, \texttt{Lampson} \longmapsto \texttt{k}_2\}$, $c(\texttt{k}_1) = \{\texttt{Ron} \longmapsto \texttt{Rivest}\}$, $c(\texttt{k}_2) = \{\texttt{Rivest} \longmapsto \texttt{k}_3\}$). Then it is easy to see that $w, \texttt{k} \not\models \texttt{k}$'s ($\texttt{Lampson}$'s $\texttt{Ron}$) $\longmapsto \texttt{k}_3$, since $[\![\texttt{k}\text{'s }(\texttt{Lampson}\text{'s }\texttt{Ron})]\!]_{w,l_w,\texttt{k}} = \emptyset$ whereas $[\![\texttt{k}_3]\!]_{w,l_w,\texttt{k}} = \{\texttt{k}_3\}$.

5 Logic Programming Implementations of Name Resolution Queries

The reader familiar with the theory of logic programming may have noted a close resemblance of the results and constructions of the preceding sections to the (now standard) fixpoint semantics for logic programs developed originally by van Emden and Kowalski [EK76]. Indeed, it is possible to translate our semantics into the framework of logic programming. In fact, we provide a translation that does not require the use of function symbols and thus produces a *Datalog* program, a restricted type of logic program that has significant computational advantages over unrestricted logic programs. Our translation allows us to take advantage of the significant body of research on the optimization of Datalog programs [Ull88, Ull89].

The idea is to translate queries to formulas in a first-order language over a vocabulary V which consists of a constant symbol for each element in $K \cup G \cup N$ and a ternary predicate symbol $\texttt{name}$. Intuitively, $\texttt{name}(x, y, z)$ says that, in the local name space of key x, the basic principal expression (i.e., key, global name or local name) y is bound to key z.

Using $\texttt{name}$, for each principal expression $\texttt{p}$ and pair of variables x, y, we define a first-order formula $\tau_{x,y}(\texttt{p})$ that, intuitively, corresponds to the assertion "$y \in [\![\texttt{p}]\!]_x$," by induction on the structure of $\texttt{p}$:

1. $\tau_{x,y}(\texttt{p}) = \texttt{name}(x, \texttt{p}, y)$ when $\texttt{p} \in K \cup G \cup N$.
2. $\tau_{x,y}(\texttt{q}\text{'s }\texttt{r}) = \exists z(\tau_{x,z}(\texttt{q}) \wedge \tau_{z,y}(\texttt{r}))$, where $z \neq x, y$.

Recall that a *Herbrand structure* over the vocabulary V is a first-order structure that has as its domain the set of constant symbols $K \cup G \cup N$ in V and interprets each constant sybol as itself. Such a structure may be represented as a set of tuples of the form $\texttt{name}(x, y, z)$, where $x, y, z \in K \cup G \cup N$. The subset relation on such sets partially orders the Herbrand structures.

We say that a Herbrand structure M over V *represents* a world $w = (\beta, c)$ and local name assignment l if, for all $x, y, z \in K \cup G \cup N$, we have $\texttt{name}(x, y, z) \in M$ iff either

1. $x, y, z \in K$ and $z = y$, or

2. $x \in K$, $y \in G$ and $z \in \beta(y)$, or

3. $x \in K$, $y \in N$ and $z \in l(x, y)$.

Intuitively, M represents w and l if it encodes all the interpretations of basic principal expressions given by w and l. The following result, whose straightforward proof is left to the reader, shows that in this case M also captures the interpretation of all other principal expressions, and expresses the correctness of our translation of principal expressions.

Proposition 5.1: *If M represents w and l then, for all principal expressions $\mathtt{p}$ and $x, y \in K \cup G \cup N$, we have $M \models \tau_{x,y}(\mathtt{p})$ iff $x, y \in K$ and $w, l, x \models \mathtt{p} \longmapsto y$.*

We now show how a logic program can be used to capture the relationsip between w and l_w. For each world $w = (\beta, c)$, we define a theory (set of sentences) Σ_w that characterizes w; Σ_w consists of the following sentences:

1. a sentence $\mathtt{name}(\mathtt{k}_1, \mathtt{k}_2, \mathtt{k}_2)$, for each pair of keys $\mathtt{k}_1, \mathtt{k}_2 \in K$, and

2. the sentence $\mathtt{name}(\mathtt{k}_1, \mathtt{g}, \mathtt{k}_2)$, for each pair of keys $\mathtt{k}_1, \mathtt{k}_2 \in K$ and global name $\mathtt{g} \in G$ such that $\mathtt{k}_2 \in \beta(\mathtt{g})$,

3. the sentence $\forall y(\tau_{\mathtt{k},y}(\mathtt{q}) \Rightarrow \mathtt{name}(\mathtt{k}, \mathtt{n}, y))$, for each key $\mathtt{k}$ and binding $\mathtt{n} \longmapsto \mathtt{q}$ in $c(\mathtt{k})$.

After some equivalence-preserving syntactic transformations (moving the existentials in the body of these sentences to the front), the theory Σ_w is a *definite Horn theory*, i.e., it consists of formulas of the form $\forall \mathbf{x}(B \Rightarrow H)$, where B is a (possibly empty) conjunction of *atoms* (that is, formulas of the form $\mathtt{name}(x, y, z)$ or $y = z$) and H is an atom. Well-known results from the theory of logic programming show that such a theory Σ has a Herbrand model M_Σ minimal with respect to the containment ordering on Herbrand structures. Moreover, this minimal Herbrand model captures the minimal name assignments for w.

Theorem 5.2: *The minimal Herbrand model M_w of Σ_w represents w and l_w.*

Using Proposition 5.1, we immediately obtain the following corollary.

Corollary 5.3: *For all $x, y \in K \cup G \cup N$ and principal expressions $\mathtt{p}$, we have $M_w \models \tau_{x,y}(\mathtt{p})$ iff $x, y \in K$ and $w, x \models_c \mathtt{p} \longmapsto y$.*

Because Σ_w is a definite Horn theory, it corresponds to a logic program. Moreover, for *existential queries*, i.e., queries ϕ that are sentences formed from atomic formulas using only conjunction, disjunction and existential quantification (but not negation), we have that Σ entails ϕ iff $M_\Sigma \models \phi$. This enables us to exploit logic programming technology to obtain efficient implementations of several types of queries, corresponding to different choices of bound and free variables in the predicate "$\mathtt{name}$". We may even form complex queries not corresponding in any direct way to the capacities of the procedure REF2. Examples of this include the following:

1. the query $\mathtt{name}(\mathtt{k}_1, \mathtt{n}, \mathtt{k}_2)$ returns "yes" if $\mathtt{k}_2$ is bound to the local name $\mathtt{n}$ according to $\mathtt{k}_1$;

2. the query $\mathtt{name}(X, \mathtt{n}, \mathtt{k})$ returns the set of keys X such that $\mathtt{k}$ is in $\mathtt{n}$ according to X;
3. the query $\mathtt{name}(\mathtt{k}_1, X, \mathtt{k}_2)$ returns the set of global and local names X containing $\mathtt{k}_2$ according to $\mathtt{k}_1$.
4. the query $\mathtt{name}(\mathtt{k}_1, \mathtt{n}, X) \wedge \mathtt{name}(\mathtt{k}_2, \mathtt{n}, X)$ returns the set of keys X that $\mathtt{k}_1$ and $\mathtt{k}_2$ agree to be associated with local name $\mathtt{n}$.

Many more possibilities clearly exist. These observations show the advantage of viewing name resolution in a logic programming framework.

6 Self

Abadi considers an extension of his logic obtained by adding a special basic principal expression `Self`, intended to represent SDSI's expression `(ref:)`. (We remark that `Self` is essentially the same as I in the logic of naming considered in [GH93].) Intuitively, `Self` denotes the current principal. The semantics given to `Self` by Abadi extends the definition of the set of principals associated with a principal expression by taking $[\![\mathtt{Self}]\!]_a = \{a\}$ for each $a \in \mathcal{W}$. This suffices to validate the following axiom.

$$\text{Identity:} \quad \begin{array}{ll} \mathtt{Self}\text{'s } \mathtt{p} \longmapsto \mathtt{p} & \mathtt{p} \longmapsto \mathtt{Self}\text{'s } \mathtt{p} \\ \mathtt{p}\text{'s } \mathtt{Self} \longmapsto \mathtt{p} & \mathtt{p} \longmapsto \mathtt{p}\text{'s } \mathtt{Self} \end{array}$$

These axioms very reasonably capture the intuitions that `Self` refers to the current principal.

However, not all consequences of this semantics for `Self` are so reasonable. For example, the following is valid under Abadi's semantics:

$$\begin{array}{l} (\mathtt{k}_P \textit{ says } \mathtt{US} \longmapsto \mathtt{Self}) \wedge (\mathtt{k}_P \textit{ says } \mathtt{US} \longmapsto \mathtt{k}_{VP}) \\ \Rightarrow \mathtt{k}_P \textit{ says } ((\mathtt{US} \textit{ says false}) \vee (\mathtt{Self} \longmapsto \mathtt{k}_{VP})) \end{array} \tag{1}$$

Interpreting $\mathtt{k}_P$ as the key of the president of the US and $\mathtt{k}_{VP}$ the key of the vice-president, this is clearly unreasonable. It should not follow from the fact the the president says that both he and the vice-president speak for the US that according to the president, either the US speaks nonsense or the vice president speaks for the president.

Abadi's suggested semantics for `Self` works much better in the context of the logic LLNC. Suppose we extend this logic to include `Self`, and like Abadi, define $[\![\mathtt{Self}]\!]_{\mathtt{k}} = \{\mathtt{k}\}$ for keys $\mathtt{k} \in K$. This again validates the Identity axioms above. To get completeness, we just need to add one axiom in addition to Identity, which basically says that `Self` acts like a key (cf. Nonemptiness (d)):

$$\text{Self-is-key} \quad \mathtt{Self} \longmapsto \mathtt{p} \wedge \mathtt{p}\text{'s } \mathtt{k} \longmapsto \mathtt{k} \Rightarrow \mathtt{p} \longmapsto \mathtt{Self}.$$

Let $\mathrm{AX}_{\mathit{inf}}^{\mathit{self}}$ (resp., $\mathrm{AX}_{\mathit{fin}}^{\mathit{self}}$) be the result of adding Identity and Self-is-key to $\mathrm{AX}_{\mathit{inf}}$ (resp., $\mathrm{AX}_{\mathit{fin}}$). Let LLNC^s be the language that results when we add `Self` to the syntax.

Theorem 6.1: *AX_{inf}^{self} (resp., AX_{fin}^{self}) is a sound and complete axiomatization of $LLNC^s$ with respect to the open semantics if K is infinite (resp., K is finite).*

Propositions 3.3 and Theorem 3.5 hold with essentially no change in proof for LLNC^s; it follows that AX^{self}_{fin} (resp., AX^{self}_{inf}) is also complete with respect to the closed semantics and the satisfiability problem is NP-complete.

Interestingly, the proof of completeness shows that once we add Identity and Self-is-key to the axioms, we no longer need Current Principal as an axiom in the finite case. Here is a sketch of the argument: From Identity we get that $\mathtt{Self}$'s $\mathtt{k} \longmapsto \mathtt{k}$ is provable for any key $\mathtt{k}$. Now applying Witnesses, we get that $\vee_{\mathtt{k} \in K} \mathtt{Self} \longmapsto \mathtt{k}$ is provable. Together with Self-is-key, this says that $\mathtt{Self}$ is one of the keys in K. Identity (together with Transitivity) tells us that for that key $\mathtt{k}$ that is Self, $\mathtt{n} \longmapsto \mathtt{k}' \Leftrightarrow \mathtt{k}$'s $\mathtt{n} \longmapsto \mathtt{k}'$ holds, giving us Current Principal.

Note that with our semantics for $\mathtt{Self}$, the counterintuitive conclusion (1) does not follow. From $\mathtt{k}_P$ *cert* $\mathtt{US} \longmapsto \mathtt{Self}$ and $\mathtt{k}_P$ *cert* $\mathtt{US} \longmapsto \mathtt{k}_{VP}$ it follows that $[\![\mathtt{US}]\!]_{\mathtt{k}_P} \supseteq \{\mathtt{k}_P, \mathtt{k}_{VP}\}$. Thus, we have neither $[\![\mathtt{US}]\!]_{\mathtt{k}_P} = \emptyset$ nor $\{\mathtt{k}_{VP}\} \supseteq [\![\mathtt{US}]\!]_{\mathtt{k}_P}$, which would be required to get a conclusion similar to that drawn by Abadi's logic.

7 Conclusions

We have introduced a logic LLNC for reasoning about SDSI's local name spaces and have argued that it has some significant advantages over Abadi's logic. Among other things, it provides a complete characterization of SDSI's REF2, has an elegant complete axiomatization, and its connections with Logic Programming lead to efficient implementations of many queries of interest.

We believe that some of the dimensions in which Abadi's logic differs from SDSI warrant further investigation. For example, under some sensible interpretations, the conclusions reached by Abadi's logic in Example 2.4 are quite reasonable. One such interpretation is that while local names may be bound to more than one key, they are intended to denote a single individual. If $\mathtt{k}$ knows that $\mathtt{k}_1$ and $\mathtt{k}_2$ are two keys used by the one individual Lampson, and Lampson uses $\mathtt{k}_1$ to certify that his local name $\mathtt{Ron}$ is bound to the name $\mathtt{Rivest}$, and also uses his key $\mathtt{k}_2$ to certify that his local name $\mathtt{Rivest}$ is bound to $\mathtt{k}_3$, then it is very reasonable to conclude that $\mathtt{k}$'s $\mathtt{Lampson}$'s $\mathtt{Ron}$ is bound to $\mathtt{k}_3$. Another interpretation supporting this conclusion would be that *says* aggregates the certificates issued using a number of distinct keys (possibly belonging to distinct individuals) much in the way that the notion of *distributed knowledge* [FHMV95] from the literature on reasoning about knowledge aggregates the knowledge of a collection of agents. We believe that our semantic framework, which, unlike Abadi's, makes the set of certificates issued explicit, provides an appropriate basis for the study of such issues.

Our semantic framework also lends itself to a number of generalizations, which we are currently exploring. These include reasoning about the beliefs of principals and reasoning about permission, authority, and delegation. We hope to report on this work shortly.

A Proofs

In this appendix, we prove all the technical results stated in the main text. For ease of exposition, we repeat the statements of the results here.

Theorem 3.1: *Given a world w, there exists a unique local name assignment l_w minimal in the set of all local name assignments consistent with w. Moreover, if $\mathtt{p}$ is a principal expression and $\mathtt{k}_1$ and $\mathtt{k}_2$ are keys, then $w, l_w, \mathtt{k}_1 \models_o \mathtt{p} \longmapsto \mathtt{k}_2$ iff, for all local name assignments l consistent with w, we have $w, l, \mathtt{k}_1 \models_o \mathtt{p} \longmapsto \mathtt{k}_2$.*

Proof: This result can be established using standard results from the theory of fixed points. Suppose $(X, \leq)$ is a complete partial order. Denote the least upper bound of a set $Y \subseteq X$ by $\sqcup Y$. A mapping $T : X \to X$ is said to be *monotonic* if for all $x \leq y$ in X we have $T(x) \leq T(y)$. Such a mapping T is said to be *continuous* if for all infinite increasing sequences $x_0 \leq x_1 \leq \ldots$ in X we have $T(\sqcup\{x_i \ : \ i \in \mathbf{N}\}) = \sqcup\{T(x_i) \ : \ i \in \mathbf{N}\}$. Note that continuity implies monotonicity. To establish continuity of a monotonic mapping T, it suffices to show that $T(\sqcup\{x_i \ : \ i \in \mathbf{N}\}) \leq \sqcup\{T(x_i) \ : \ i \in \mathbf{N}\}$, since the opposite containment is immediate from monotonicity.

For a fixed expression $\mathtt{p}$, world w and key $\mathtt{k}$, the expression $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}}$ is easily seen to be monotonic in l, i.e., if $l \leq l'$ then $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}} \subseteq [\![\mathtt{p}]\!]_{w,l',\mathtt{k}}$. Moreover, it is also continuous in l.

Lemma A.1: *Suppose $l_0 \leq l_1 \leq \ldots$ is an increasing sequence of local name assignments and let $l_\omega = \sqcup_{m \in \mathbf{N}} l_m$. For all principal expressions $\mathtt{p}$, we have $[\![\mathtt{p}]\!]_{w,l_\omega,\mathtt{k}} = \bigcup_{m \in \mathbf{N}} [\![\mathtt{p}]\!]_{w,l_m,\mathtt{k}}$.*

Proof: By a straightforward induction on the structure of $\mathtt{p}$. ∎

Given the world $w = (\beta, c)$, we define an operator T_w on the space of local name assignments *LNA*. For a local name assignment l, we define $T_w(l)$ to be the local name assignment such that for all $\mathtt{k} \in K$ and $\mathtt{n} \in N$, the set $T_w(l)(\mathtt{k}, \mathtt{n})$ is the union of the sets $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}}$ such that the formula $\mathtt{n} \longmapsto \mathtt{p}$ is in $c(\mathtt{k})$. The following lemma is follows easily from Lemma A.1.

Lemma A.2: *The mapping T_w is a continuous operator on $(LNA, \leq)$.*

The following lemma is almost immediate from the definitions.

Lemma A.3: *A local name assignment l is consistent with a world w iff $T_w(l) \leq l$.*

Suppose $(X, \leq)$ is a complete partial order with minimal element $\bot$. An element $x \in X$ is said to be a *pre-fixpoint* of an operator T on X if $T(x) \leq x$; x is a *fixpoint* of T if $T(x) = x$. Given an operator T on X, define a sequence of elements $T \uparrow \gamma$, where γ is an ordinal, as follows. For the base case, let $T \uparrow 0 = \bot$. For successor ordinals $\gamma + 1$, define $T \uparrow \gamma + 1 = T(T \uparrow \gamma)$. For limit ordinals γ, define $T \uparrow \gamma = \sqcup\{T \uparrow \delta \ : \ \delta < \gamma\}$. A well-known result (see [LNS82] for a discussion of its history) states that if T is continuous then then this sequences converges to the least pre-fixpoint of T, that convergence has taken place by $\gamma = \omega$, and that $T \uparrow \omega$ is in fact a fixed point of T. Thus, we obtain as a corollary of Lemma A.2 and Lemma A.3 that there exists a minimal local name assignment consistent with w, and that this local name assignment equals $T_w \uparrow \omega$. The second half of Theorem 3.1 is immediate from the earlier observation that $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}}$ is monotonic in l. ∎

Theorem 3.2: *AX_{fin} (resp., AX_{inf}) is a sound and complete axiomatization of LLNC with respect to the open semantics if K is infinite (resp., K is finite).*

Proof: We start with the completeness proof for AX_{inf}, so that we assume that K is infinite. We then show how to deal with AX_{fin}. As usual, it suffices to show that if ϕ is AX_{inf}-consistent, then ϕ is satisfable. In fact, we put a little extra work into our proof that ϕ is satisfiable so that we can prove Proposition 3.3 as well.

Let $\mathrm{Sub}(\phi)$ consist of all subformulas of ϕ. We say that a principal expression $\mathtt{p}'$ is a *variant* of $\mathtt{p}$ if $\mathtt{p} \longmapsto \mathtt{p}'$ and $\mathtt{p}' \longmapsto \mathtt{p}$ are both provable using only Reflexivity, Associativity, and Transitivity. The *left-associative* variant of a principal expression $\mathtt{p}$ is the one where we associate all terms to the left. Thus, $((\mathtt{n}_1\text{'s } \mathtt{n}_2)\text{'s } \mathtt{n}_3)\text{'s } \mathtt{n}_4$ is the left-associative variant of $\mathtt{n}_1\text{'s } ((\mathtt{n}_2\text{'s } \mathtt{n}_3)\text{'s } \mathtt{n}_4)$.

Define P to be the smallest set of principal expressions such that

1. if $\mathtt{p} \longmapsto \mathtt{q}$ is in $\mathrm{Sub}(\phi)$ then $\mathtt{p}$ and $\mathtt{q}$ are in P,
2. if $\mathtt{k}$ *cert* $(\mathtt{n} \longmapsto \mathtt{p}) \in \mathrm{Sub}(\phi)$ then $\mathtt{k}\text{'s } \mathtt{n}$ and $\mathtt{k}\text{'s } \mathtt{p}$ are in P,
3. if $\mathtt{p} \in P$ and $\mathtt{p}'$ is the left-associative variant of $\mathtt{p}$, then $\mathtt{p}' \in P$,
4. P is closed under subexpressions, so that if $\mathtt{p}\text{'s } \mathtt{q} \in P$, then so are $\mathtt{p}$ and $\mathtt{q}$,
5. if $\mathtt{k} \in P$ is a key and $\mathtt{n} \in P$ is a local name, then $\mathtt{k}\text{'s } \mathtt{n} \in P$.

For Proposition 3.3, it is necessary to get an upper bound on the size of P in terms of $|\phi|$.

Lemma A.4: $|P| < 2 \cdot |\phi|^2$.

Proof: Let $|\mathtt{p}|$ be the total number of expressions in $G \cup K \cup N$ that appear in $\mathtt{p}$, counted with multiplicity. An easy proof by induction on structure shows that a principal expression $\mathtt{p}$ has at most $|\mathtt{p}|$ subexpressions, at least one of which must be in $G \cup K \cup N$. For every other subexpression $\mathtt{q}$, there is a unique left-associative variant $\mathtt{q}'$, which has at most $|\mathtt{q}'| = |\mathtt{q}| \leq |\mathtt{p}|$ subexpressions, each of which is associated to the left. Thus, starting with a principal expression $\mathtt{p}$, the least set closed under clauses 3 and 4 above contains at most $|\mathtt{p}|^2$ elements. Now a straightforward induction on the structure of ϕ shows that the least set P' closed under clauses 1-4 above has at most $|\phi|^2$ expressions. Finally, it is easy to see that closing off under 5 gives us P, since the set that results after closing off under 5 is still closed under 1–4. Moreover, this final step adds at most $|\phi|^2$ expressions $\mathtt{k}\text{'s } \mathtt{n}$, since both $\mathtt{k}$ and $\mathtt{n}$ must be subexpressions of ϕ. ∎

Let $\mathtt{k}_0$ be some key not occurring in P. We use $\mathtt{k}_0$ both to express emptiness of expressions in P and as the "current principal". Define P_1 to be the set of principal expressions $P \cup \{\mathtt{k}_0\} \cup \{\mathtt{p}\text{'s } \mathtt{k}_0 \ : \ \mathtt{p} \in P\}$. Let E be consist of the formulas $\mathtt{p}\text{'s } \mathtt{k}_0 \longmapsto \mathtt{k}_0$ for each $\mathtt{p} \in P$. Note that all principal expressions occurring in the formulas in E are in P_1. Let S be an AX_{inf}-consistent set containing ϕ and, for every formula $\psi \in \mathrm{Sub}(\phi) \cup E$, either ψ or $\neg\psi$. Since ϕ is AX_{inf}-consistent, there must be some AX_{inf}-consistent set S of this form.

Define $S^+ = Cl(S, P_1)$ to be the smallest set of formulas containing S closed under Reflexivity, Transitivity, Left Motonocity, Converse of Globality, Globality, and Nonemptiness, in the sense that

(ClR) if $\mathtt{p} \in P_1$, then $\mathtt{p} \longmapsto \mathtt{p} \in S^+$,

(ClT) if $\mathtt{p} \longmapsto \mathtt{q}$ and $\mathtt{q} \longmapsto \mathtt{r}$ are both in S^+, then $\mathtt{p} \longmapsto \mathtt{r} \in S^+$,

(ClLM) if $\mathtt{p} \longmapsto \mathtt{q} \in S^+$, $\mathtt{p}$'s $\mathtt{r} \in P_1$, and $\mathtt{q}$'s $\mathtt{r} \in P_1$, then $\mathtt{p}$'s $\mathtt{r} \longmapsto \mathtt{q}$'s $\mathtt{r} \in S^+$,

(ClCG) if $\mathtt{p}$'s $\mathtt{g} \in P_1$ for $\mathtt{g} \in K \cup G$ then $\mathtt{g} \longmapsto \mathtt{p}$'s $\mathtt{g} \in S^+$,

(ClG) if $\mathtt{p}$'s $\mathtt{k} \longmapsto \mathtt{k} \in S^+$ for some key $\mathtt{k}$ and $\mathtt{p}$'s $\mathtt{g} \in P_1$, where $\mathtt{g} \in K \cup G$, then $\mathtt{p}$'s $\mathtt{g} \longmapsto \mathtt{g} \in S^+$,

(ClKL) if $\mathtt{k}$ *cert* $(\mathtt{n} \longmapsto \mathtt{p}) \in S^+$ then $(\mathtt{k}$'s $\mathtt{n} \longmapsto \mathtt{k}$'s $\mathtt{p}) \in S^+$,

(ClK) if $\mathtt{p} \longmapsto \mathtt{k}' \in S^+$ and $\mathtt{p}$'s $\mathtt{k} \in P_1$, then $\mathtt{p}$'s $\mathtt{k} \longmapsto \mathtt{k} \in S^+$,

(ClN) if $\neg(\mathtt{p} \longmapsto \mathtt{q}) \in S^+$ and $\mathtt{q}$'s $\mathtt{k} \in P_1$, then $\mathtt{q}$'s $\mathtt{k} \longmapsto \mathtt{k} \in S^+$,

(ClC) if $\mathtt{p}$'s $\mathtt{q} \longmapsto \mathtt{k}_1 \in S^+$ and $\mathtt{p}$'s $\mathtt{k} \in P_1$, then $\mathtt{p}$'s $\mathtt{k} \longmapsto \mathtt{k} \in S^+$,

(ClNE) if $\mathtt{p}$'s $\mathtt{k} \longmapsto \mathtt{k}$ and $\mathtt{k}' \longmapsto \mathtt{p}$ are both in S^+, then $\mathtt{p} \longmapsto \mathtt{k}' \in S^+$,

(ClKD) if $\mathtt{k}$ and $\mathtt{k}'$ are distinct keys in P, then $\neg(\mathtt{k} \longmapsto \mathtt{k}') \in S^+$,

(ClLV) If $\mathtt{p}'$ is the left-associative variant of $\mathtt{p} \in P$, then $\mathtt{p} \longmapsto \mathtt{p}' \in S^+$ and $\mathtt{p}' \longmapsto \mathtt{p} \in S^+$.

It is easy to see that S^+ is AX_{inf}-consistent, since S is and each of the closure rules emulates an axiom in AX_{inf}. Our goal now is to show that there exists a triple $w, l, \mathtt{k}$ such that $w, l, \mathtt{k} \models \psi$ for all $\psi \in S$ (and thus, in particular, $w, l, \mathtt{k} \models \phi$).

Lemma A.5: *If* $\mathtt{k}_0$ *appears in the formula* $\mathtt{p} \longmapsto \mathtt{q} \in S^+$*, then* $\mathtt{k}_0$ *appears in both* $\mathtt{p}$ *and* $\mathtt{q}$.

Proof: An easy induction on the construction of S^+, using the fact that all principal expressions occurring in S^+ are in P_1 and $\mathtt{k}_0$ appears only as the right most expression in a principal expression in P_1. ∎

By Lemma A.5, if $\mathtt{p} \longmapsto \mathtt{q} \in S^+$ and one of the expressions $\mathtt{p}, \mathtt{q}$ is in P (and thus does not mention $\mathtt{k}_0$) then so is the other. Define a binary relation $\approx$ on P by defining $p \approx q$ if both $\mathtt{p} \longmapsto \mathtt{q}$ and $\mathtt{q} \longmapsto \mathtt{p}$ are in S^+. It is immediate from transitivity and reflexivity that $\approx$ is an equivalence relation on P. Given $\mathtt{p} \in P$, we write $[\mathtt{p}]$ for the equivalence class of $\mathtt{p}$ under $\approx$.

We classify the expressions in P as follows. Say that an expression $\mathtt{p}$ in P is *empty* (with respect to S^+) if $\neg(\mathtt{p}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0)$ is in S^+. Say that $\mathtt{p}$ is *key-equivalent* if it is not empty and $\mathtt{k} \longmapsto \mathtt{p}$ is in S^+ for some key $\mathtt{k}$ (by (ClNE) this implies $\mathtt{p} \approx \mathtt{k}$). Intuitively, the interpretation of an empty expression will be the empty set and the interpretation of a key-equivalent expression $\mathtt{p}$ such that $\mathtt{k} \longmapsto \mathtt{p} \in S^+$ will be $\{\mathtt{k}\}$. If $\mathtt{p}$ is neither empty nor key-equivalent, we say it is *open*. Clearly, every expression in P is either empty, key-equivalent, or open. Moreover, by (ClLM) and (ClT), if $\mathtt{p} \approx \mathtt{q}$ then $\mathtt{p}$ is empty, key-equivalent or open iff $\mathtt{q}$ is. In particular, we may sensibly refer to open $\approx$-equivalence classes of expressions in P.

Let O be the set of open equivalence classes of expressions in P. Note that if $K_\phi \subseteq K$ consists of all the keys in K that appear in ϕ, then there are fewer than $2 \cdot |\phi|^2 - |K_\phi|$ equivalence classes of open expressions. For each class $c \in O$, let $\mathtt{k}_c$ be a fresh key. Intuitively, the key $\mathtt{k}_c$ will act as a canonical representative of the keys in the interpretation of an expression $\mathtt{p} \in c$, in the sense that the interpretations of $\mathtt{p}$'s $\mathtt{q}$ and $\mathtt{k}_c$'s $\mathtt{q}$ will be the same for certain expressions $\mathtt{q}$. Since K is infinite, we are guaranteed that we can always find keys $\mathtt{k}_c$, but the argument works even if K is finite, as long as $|K| \geq 2 \cdot |\phi|^2$. (We also need to have a key in $K \setminus K_\phi$ to be $\mathtt{k}_0$.)

Define S^* to be consist of S^+ together with, for all $c \in O$,

1. the formula $\mathtt{k}_c \longmapsto \mathtt{k}_c$, and

2. the formulas $\mathtt{p} \longmapsto \mathtt{k}_c$, where for some $\mathtt{q} \in c$ we have $\mathtt{p} \longmapsto \mathtt{q} \in S^+$.

It is easy to show that $\mathtt{k}_0$ does not appear in any formula in $S^* - S^+$: Clearly $\mathtt{k}_0$ does not appear in the formulas $\mathtt{k}_c \longmapsto \mathtt{k}_c$ added by clause 1. If $\mathtt{p} \longmapsto \mathtt{q}$ is a formula added by clause 2, then there is some equivalence class c and expression $\mathtt{q} \in c$ such that $\mathtt{p} \longmapsto \mathtt{q} \in S^+$. Since c is an equivalence class of expressions in P, none of which contain $\mathtt{k}_0$, the expression $\mathtt{q}$ does not contain $\mathtt{k}_0$. It follows from Lemma A.5 that $\mathtt{p}$ does not contain $\mathtt{k}_0$. Since $S^* - S^+$ contains no formulas involving $\mathtt{k}_0$, S^* also satisfies the property stated for S^+ in Lemma A.5.

Define the local name assignment l as follows. Given a key $\mathtt{k}$ and local name $\mathtt{n}$,

1. $l(\mathtt{k}_0, \mathtt{n}) = \{\mathtt{k}' \in K \mid \mathtt{n} \longmapsto \mathtt{k}' \in S^*\}$,

2. $l(\mathtt{k}, \mathtt{n}) = \{\mathtt{k}' \in K \mid \mathtt{k}\text{'s } \mathtt{n} \longmapsto \mathtt{k}' \in S^*\}$ if $\mathtt{k} \in P$,

3. $l(\mathtt{k}, \mathtt{n}) = \{\mathtt{k}' \in K \mid \mathtt{p}\text{'s } \mathtt{n} \longmapsto \mathtt{k}' \in S^* \text{ and } \mathtt{p} \in c\}$ if $\mathtt{k} = \mathtt{k}_c$ for some $c \in O$,

4. $l(\mathtt{k}, \mathtt{n}) = \emptyset$ for all other $\mathtt{k}$.

Define the world $w = (\beta, c)$ by taking $\beta(\mathtt{g}) = \{\mathtt{k} \in K \mid \mathtt{g} \longmapsto \mathtt{k} \in S^*\}$ and defining $c(\mathtt{k})$, for each key $\mathtt{k}$, to be the set of formulas $\mathtt{n} \longmapsto \mathtt{p}$ such that $(\mathtt{k}\ \mathit{cert}\ (\mathtt{n} \longmapsto \mathtt{p})) \in S$. Note for future reference that there exists a finite subset K_1 of K such that $l(\mathtt{n}, \mathtt{k}) \subseteq K_1$, $l(\mathtt{n}, \mathtt{k}) = \emptyset$ for $\mathtt{k} \notin K_1$, $\beta(\mathtt{g}) \subseteq K_1$, and $\beta(\mathtt{g}) = \emptyset$ if $\mathtt{g}$ does not appear in ϕ. Indeed, K_1 consists of the keys that appear in S, $\mathtt{k}_0$, and the keys $\mathtt{k}_c$ for $c \in O$.

Let $I(\mathtt{p}) = \{\mathtt{k} \in K \mid \mathtt{p} \longmapsto \mathtt{k} \in S^*\}$.

Lemma A.6: *If* $\mathtt{p} \in P$, *then* $\mathtt{p}$ *is empty iff* $I(\mathtt{p}) = \emptyset$.

Proof: If $\mathtt{p}$ is not empty, then it is either key-equivalent or open. If it is key-equivalent, we have already observed that there must exist some key $\mathtt{k}'$ such that $\mathtt{p} \longmapsto \mathtt{k}' \in S^*$, so $I(\mathtt{p}) \neq \emptyset$. If it is open, suppose it is in equivalence class c. Then $\mathtt{p} \longmapsto \mathtt{k}_c \in S^*$, since $\mathtt{p} \longmapsto \mathtt{p} \in S^+$ by (ClR). Again, it follows that $I(\mathtt{p}) \neq \emptyset$.

Conversely, suppose that $I(\mathtt{p}) \neq \emptyset$. Thus, $\mathtt{p} \longmapsto \mathtt{k} \in S^*$ for some key $\mathtt{k}$. If $\mathtt{p} \longmapsto \mathtt{k} \in S^+$, then by (ClK), $\mathtt{p}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0 \in S^+$, so $\mathtt{p}$ is not empty. If $\mathtt{p} \longmapsto \mathtt{k} \notin S^+$, then $\mathtt{k} = \mathtt{k}_c$, and there is some $\mathtt{q} \in c$ such that $\mathtt{p} \longmapsto \mathtt{q} \in S^+$. Since $\mathtt{q}$ is open, $\mathtt{q}$ cannot be empty, so $\mathtt{q}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0 \in S^+$. Moreover, by (ClLM), $\mathtt{p}$'s $\mathtt{k}_0 \longmapsto \mathtt{q}$'s $\mathtt{k}_0 \in S^+$. Thus, by (ClT), $\mathtt{p}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0 \in S^+$, so $\mathtt{p}$ is nonempty. ∎

Lemma A.7: *For all expressions* $\mathtt{p} \in P$, *we have* $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}_0} = I(\mathtt{p})$.

Proof: We proceed by induction on $|\mathtt{p}|$ (as defined in Lemma A.4). The claim is immediate from the definitions in case $\mathtt{p}$ is a global name or a local name. Suppose that $\mathtt{p}$ is a key $\mathtt{k}_1$. Then $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}_0} = \{\mathtt{k}_1\}$. Since $\mathtt{k}_1 \longmapsto \mathtt{k}_1 \in S^*$ by construction, it follows that $\mathtt{k}_1 \in I(\mathtt{k}_1)$. It remains to show that $I(\mathtt{k}_1) \subseteq \{\mathtt{k}_1\}$. Suppose $(\mathtt{k}_1 \longmapsto \mathtt{k}) \in S^*$. By Lemma A.5, we cannot have $\mathtt{k} = \mathtt{k}_0$. Since S^+ is AX$_{inf}$-consistent and closed under (ClKD), if $\mathtt{k} \in P$ we must have $\mathtt{k}_1 = \mathtt{k}$. The remaining possibility for $\mathtt{k}$, that it equals $\mathtt{k}_c$ for some $c \in O$, cannot happen. For if so, only the second clause of the definition of S^* could explain $(\mathtt{k}_1 \longmapsto \mathtt{k}) \in S^*$. But then we have $(\mathtt{k}_1 \longmapsto \mathtt{q}) \in S^+$ for some $\mathtt{q} \in c$. This contradicts the assumption that c is an equivalence class of open expressions.

Finally, suppose that $|\mathtt{p}| > 1$. Let $\mathtt{p}'$ be the left-associative variant of $\mathtt{p}$. It is clear from the semantics that $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}_0} = [\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0}$. Morover, (ClLV) and (ClT) guarantee that $I(\mathtt{p}) = I(\mathtt{p}')$. Thus, it suffices to prove that $I(\mathtt{p}') = [\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0}$. Suppose that $\mathtt{p}' = \mathtt{q}$'s $\mathtt{r}$. The definition of length guarantees that $|\mathtt{p}'| = |\mathtt{p}| > |\mathtt{q}|$, so the induction hypothesis applies to $\mathtt{q}$. Since $\mathtt{p}'$ is associated to the left, $\mathtt{r} \in G \cup K \cup N$.

Suppose that $\mathtt{r} = \mathtt{g} \in G \cup K$. Note that $[\![\mathtt{q}\text{'s }\mathtt{g}]\!]_{w,l,\mathtt{k}_0} = \emptyset$ if $[\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0} = \emptyset$ and $[\![\mathtt{q}\text{'s }\mathtt{g}]\!]_{w,l,\mathtt{k}_0} = [\![\mathtt{g}]\!]_{w,l,\mathtt{k}_0}$ if $[\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0} \neq \emptyset$. We consider these two cases separately.

Suppose first that $[\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0} = \emptyset$, so $[\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0} = \emptyset$. By the induction hypothesis, $I(\mathtt{q}) = \emptyset$. To show that $I(\mathtt{p}') = \emptyset$, we show that $\mathtt{p}'$ is empty. Suppose not. Then $(\mathtt{p}')$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0 \in S^+$. Since S^+ contains either $\mathtt{q}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0$ or $\neg(\mathtt{q}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0)$ and S^+ is AX$_{inf}$-consistent, by Nonemptiness(c), Associativity, and Transitivity, we must have $\mathtt{q}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0 \in S^+$. Thus, $\mathtt{q}$ is not empty. By Lemma A.6, $I(\mathtt{q}) \neq \emptyset$, a contradiction. Hence, $\mathtt{p}'$ is empty. It now follows from Lemma A.6 that $I(\mathtt{p}') = \emptyset$, as desired.

Consider next the case where $[\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0} \neq \emptyset$, so $[\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0} = [\![\mathtt{q}\text{'s }\mathtt{g}]\!]_{w,l,\mathtt{k}_0} = [\![\mathtt{g}]\!]_{w,l,\mathtt{k}_0}$. To show that $[\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0} = I(\mathtt{p}')$, we show that $I(\mathtt{p}') = I(\mathtt{g})$. The result then follows from the induction hypothesis.

By the induction hypothesis, $I(\mathtt{q}) \neq \emptyset$, so by Lemma A.6, $\mathtt{q}$ is not empty. It follows from (ClG) that $\mathtt{q}$'s $\mathtt{g} \longmapsto \mathtt{g} \in S^+$. Suppose that $\mathtt{k} \in I(\mathtt{g})$. If $\mathtt{k} \in P_1$, then $\mathtt{g} \longmapsto \mathtt{k} \in S^+$, so by (ClT), $\mathtt{q}$'s $\mathtt{g} \longmapsto \mathtt{k} \in S^+$ and $\mathtt{k} \in I(\mathtt{p}')$. If $\mathtt{k} = \mathtt{k}_c$ for some $c \in O$, then $\mathtt{g} \longmapsto \mathtt{q}' \in S^+$ for some $\mathtt{q}' \in c$. Thus, $\mathtt{p}' \longmapsto \mathtt{q}' \in S^+$ by (ClT) and we obtain that $\mathtt{p}' \longmapsto \mathtt{k} \in S^*$ by construction of S^*. Thus, $I(\mathtt{g}) \subseteq I(\mathtt{p}')$.

For the opposite containment, note that by (ClCG) we have $\mathtt{g} \longmapsto \mathtt{q}$'s $\mathtt{g} \in S^+$. Arguing as above, we obtain using (ClT) that $I(\mathtt{g}) \supseteq I(\mathtt{p}')$. This completes the proof that $I(\mathtt{p}') = I(\mathtt{g})$.

It remains to deal with the case that $\mathtt{p}'$ has of the form $\mathtt{q}$'s $\mathtt{n}$, where $\mathtt{n}$ is a local name. There are three possibilities: $\mathtt{q}$ is empty, key-equivalent or open. If $\mathtt{q}$ is empty, then by Lemma A.6 and the induction hypothesis, $I(\mathtt{q}) = \emptyset$ and $[\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0} = \emptyset$. It follows that $[\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0} = \emptyset$. Moreover, using Nonemptiness(c), Associativity, and Transitivity as above, it follows that $\mathtt{p}'$ is empty and hence by Lemma A.6, $I(\mathtt{p}') = \emptyset$, as desired.

If $\mathtt{q}$ is key-equivalent, say $\mathtt{q} \approx \mathtt{k}_1$, then $\mathtt{q} \longmapsto \mathtt{k}_1 \in S^+$ and $\mathtt{k}_1 \longmapsto \mathtt{q} \in S^+$. Using Key Distinctness and the consistency of S^+, it easily follows that $I(\mathtt{q}) = \{\mathtt{k}_1\}$. By the induction hypothesis, $[\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0} = \{k_1\}$. Thus, $[\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0} = l(\mathtt{k}_1, \mathtt{n})$. By construction, $l(\mathtt{k}_1, \mathtt{n}) = I(\mathtt{k}_1\text{'s }\mathtt{n}) = I(\mathtt{p}')$, as desired.

Finally, suppose that $\mathtt{q}$ is open. If $\mathtt{k} \in I(\mathtt{p}')$, then it is immediate from the construction that that $\mathtt{q} \longmapsto \mathtt{k}_{[\mathtt{q}]} \in S^*$ and $\mathtt{k} \in l(\mathtt{k}_{[\mathtt{q}]}, \mathtt{n})$. By the induction hypothesis, $\mathtt{k}_{[\mathtt{q}]} \in [\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0}$, so

$\mathtt{k} \in [\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0} = \cup_{\mathtt{k}' \in [\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0}} l(\mathtt{k}', \mathtt{n})$. Thus, $I(\mathtt{p}') \subseteq [\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0}$ if $\mathtt{p}'$ is open.

For the opposite containment, suppose that $\mathtt{k} \in [\![\mathtt{p}']\!]_{w,l,\mathtt{k}_0}$. This means that there is some key $\mathtt{k}'$ such that $\mathtt{k}' \in [\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0}$ and $\mathtt{k} \in l(\mathtt{k}', \mathtt{n})$. By the induction hypothesis, $\mathtt{k}' \in I(\mathtt{q})$, so $\mathtt{q} \longmapsto \mathtt{k}' \in S^*$. If $\mathtt{k}' \in P_1$, then $\mathtt{q} \longmapsto \mathtt{k}' \in S^+$ and $(\mathtt{k}')$'s $\mathtt{n} \longmapsto \mathtt{k} \in S^+$. (Since $\mathtt{q} \in P$ and $\mathtt{q} \longmapsto \mathtt{k}' \in S^*$, we cannot have $\mathtt{k}' = \mathtt{k}_0$, by Lemma A.5.) By (ClLM), $\mathtt{q}$'s $\mathtt{n} \longmapsto (\mathtt{k}')$'s $\mathtt{n} \in S^+$, so by (ClT) we get $\mathtt{q}$'s $\mathtt{n} \longmapsto \mathtt{k} \in S^+$. Hence, $\mathtt{k} \in I(\mathtt{p}')$. If $\mathtt{k}' = \mathtt{k}_c$, where c is an open equivalence class, then from $\mathtt{q} \longmapsto \mathtt{k}' \in S^*$ it follows that $\mathtt{q} \longmapsto \mathtt{q}' \in S^+$ for some $\mathtt{q}' \in c$. From $\mathtt{k} \in l(\mathtt{k}_c, \mathtt{n})$ it follows that $(\mathtt{r}')$'s $\mathtt{n} \longmapsto \mathtt{k} \in S^*$ for some $\mathtt{r}' \in c$. By construction of S^+ we must have $(\mathtt{r}')$'s $\mathtt{n} \in P_1$, and since $\mathtt{r}' \approx \mathtt{q}'$, we have $\mathtt{q}' \longmapsto \mathtt{r}' \in S^+$. By (ClT) we obtain $\mathtt{q} \longmapsto \mathtt{r}' \in S^+$, and hence by (ClLM) that $\mathtt{q}$'s $\mathtt{n} \longmapsto (\mathtt{r}')$'s $\mathtt{n} \in S^+$. Now notice that it follows from $\mathtt{q}$'s $\mathtt{n} \longmapsto (\mathtt{r}')$'s $\mathtt{n} \in S^+$ and $(\mathtt{r}')$'s $\mathtt{n} \longmapsto \mathtt{k} \in S^*$ that $\mathtt{q}$'s $\mathtt{n} \longmapsto \mathtt{k} \in S^*$. If $\mathtt{k} \in P_1$, this is immediate from (ClT). In case $\mathtt{k} = \mathtt{k}_d$ for some open class d, we have $(\mathtt{r}')$'s $\mathtt{n} \longmapsto \mathtt{t} \in S^+$ for some $\mathtt{t} \in d$. But then $\mathtt{q}$'s $\mathtt{n} \longmapsto \mathtt{t} \in S^+$ by (ClT); by definition of S^* we get that $\mathtt{q}$'s $\mathtt{n} \longmapsto \mathtt{k} \in S^*$. This completes the proof. ∎

Lemma A.8: *For all formulas $\psi \in Sub(\phi) \cup E$, we have $\psi \in S$ iff $w, l, \mathtt{k}_0 \models_o \psi$.*

Proof: We first show that by induction on the structure of $\psi \in \mathrm{Sub}(\phi) \cup E$ that $\psi \in S$ iff $w, l, \mathtt{k}_0 \models \psi$, and then show that the assignment l is consistent with w.

It is immediate from the construction of w that $w, l, \mathtt{k}_0 \models \psi$ iff $\psi \in S$ for ψ of the form $\mathtt{k}$ *cert* $(\mathtt{n} \longmapsto \mathtt{p})$.

If ψ has the form $\mathtt{p} \longmapsto \mathtt{q}$, note that $w, l, \mathtt{k}_0 \models \mathtt{p} \longmapsto \mathtt{q}$ iff $[\![\mathtt{p}]\!]_{w,l,\mathtt{k}_0} \supseteq [\![\mathtt{q}]\!]_{w,l,\mathtt{k}_0}$ iff (by Lemma A.7) iff $I(\mathtt{p}) \supseteq I(\mathtt{q})$. Thus, it suffices to show that $I(\mathtt{p}) \supseteq I(\mathtt{q})$ iff $\mathtt{p} \longmapsto \mathtt{q} \in S^+$, for $\mathtt{p}, \mathtt{q} \in P$.

The "if" direction is immediate from (ClT): If $\mathtt{k} \in I(\mathtt{q})$ then $\mathtt{q} \longmapsto \mathtt{k} \in S^*$, so by (ClT) and the construction of S^*, $\mathtt{p} \longmapsto \mathtt{k} \in S^*$ and thus $\mathtt{k} \in I(\mathtt{p})$.

For the "only if" direction, suppose by way of contradiction that $I(\mathtt{p}) \supseteq I(\mathtt{q})$ but $\mathtt{p} \longmapsto \mathtt{q} \notin S^+$. Then, by construction, $\neg(\mathtt{p} \longmapsto \mathtt{q}) \in S^+$. We consider three cases, depending on whether $\mathtt{q}$ is empty, key-equivalent, or open.

Note first that $\mathtt{q}$ cannot be empty: $\neg(\mathtt{p} \longmapsto \mathtt{q}) \in S^+$, so by (ClN) we have $\mathtt{q}$'s $\mathtt{k}_0 \longmapsto \mathtt{k}_0 \in S^+$.

Suppose that $\mathtt{q}$ is key-equivalent, with $\mathtt{k} \longmapsto \mathtt{q} \in S^+$. If $\mathtt{p} \longmapsto \mathtt{k} \in S^+$ then, by (ClT), $\mathtt{p} \longmapsto \mathtt{q} \in S^+$, but this is not possible because S^+ is AX_{inf}-consistent. Thus $\mathtt{p} \longmapsto \mathtt{k} \notin S^+$. Since $\mathtt{k} \in P$, $\mathtt{p} \longmapsto \mathtt{k} \notin S^*$, and thus $\mathtt{k} \in I(\mathtt{p}) - I(\mathtt{q})$, giving us the desired contradiction.

Finally, suppose $\mathtt{q}$ is open. By construction, $\mathtt{q} \longmapsto \mathtt{k}_{[\mathtt{q}]} \in S^*$. Moreover, we cannot have $\mathtt{p} \longmapsto \mathtt{k}_{[\mathtt{q}]} \in S^*$, for then there would exist $\mathtt{r} \approx \mathtt{q}$ such that $\mathtt{p} \longmapsto \mathtt{r} \in S^+$. Using (ClT), it would follow that $\mathtt{p} \longmapsto \mathtt{q} \in S^+$, which is impossible since S^+ is AX_{inf}-consistent. Thus, $\mathtt{k}_{[\mathtt{q}]} \in I(\mathtt{p}) - I(\mathtt{q})$, giving the required contradiction, and completing the proof in the case that ψ is of the form $\mathtt{p} \longmapsto \mathtt{q}$.

If ψ is of the form $\neg\psi'$ or $\psi_1 \wedge \psi_2$, the result is immediate from the induction hypothesis (in the latter case, we need the fact that if $\psi_1 \wedge \psi_2 \in \mathrm{Sub}(\phi) \cup E$, then in fact $\psi \wedge \psi_2 \in \mathrm{Sub}(\phi)$, so $\psi_1, \psi_2 \in \mathrm{Sub}(\phi)$ and the induction hypothesis applies). This completes the induction proof.

To show that the assignment l is consistent with w, suppose that $\mathtt{n} \longmapsto \mathtt{p} \in c(\mathtt{k})$. Then, by construction, $\mathtt{k}$ *cert* $(\mathtt{n} \longmapsto \mathtt{p}) \in S$. By (ClKL), we have $\mathtt{k}$'s $\mathtt{n} \longmapsto \mathtt{k}$'s $\mathtt{p} \in S^+$. By what we have just shown $w, l, \mathtt{k}_0 \models \mathtt{k}\text{'s } \mathtt{n} \longmapsto \mathtt{k}\text{'s } \mathtt{p}$. It follows that $w, l, \mathtt{k} \models \mathtt{n} \longmapsto \mathtt{p}$. Thus, l is consistent with w. ▮

Thus, we have shown that ϕ is satisfiable, completing the proof of Theorem 3.2 in the case that K is infinite. The same argument works without change if K is finite but $|K| \geq 2 \cdot |\phi|^2$. (A consequence of this is that we do not need to use the axioms Witnesses and Current Principal to derive a valid formula ϕ in $\mathrm{AX}_{\mathit{fin}}$ if $2 \cdot |\phi|^2 \leq |K|$.) Moreover, the proof shows that Proposition 3.3 holds if $|K| \geq 2 \cdot |\phi|^2$.

Now suppose that $K \leq 2 \cdot |\phi|^2$. We show that if ϕ is $\mathrm{AX}_{\mathit{fin}}$-consistent, then ϕ is satisfiable. The proof is in the spirit of that in the case of $\mathrm{AX}_{\mathit{inf}}$, but simpler.

Now let P be the least set of principal expressions containing all principal expressions that appear in ϕ and closed under subexpressions. Let F consist of all formulas of the form $\mathtt{p} \longmapsto \mathtt{k}'$ and $\mathtt{k}\text{'s } \mathtt{p} \longmapsto \mathtt{k}'$, where $\mathtt{p} \in P$ and $\mathtt{k}, \mathtt{k}' \in K$. Let S be an $\mathrm{AX}_{\mathit{fin}}$-consistent set containing ϕ and, for every formula $\psi \in \mathrm{Sub}(\phi) \cup F$, either ψ or $\neg\psi$. Since ϕ is $\mathrm{AX}_{\mathit{fin}}$-consistent, there must be some $\mathrm{AX}_{\mathit{fin}}$-consistent set S of this form.

There must be some key $\mathtt{k}_0 \in K$ such that for every local name in P and key $\mathtt{k} \in K$, we have $\mathtt{n} \longmapsto \mathtt{k} \in S$ iff $\mathtt{k}_0\text{'s } \mathtt{n} \longmapsto \mathtt{k} \in S$. For otherwise, for each key $\mathtt{k}$, there is some local name $\mathtt{n}_\mathtt{k}$ and key $\mathtt{k}_\mathtt{k}$ such that either both $\mathtt{n}_\mathtt{k} \longmapsto \mathtt{k}_\mathtt{k}$ and $\neg(\mathtt{k}\text{'s } \mathtt{n}_\mathtt{k} \longmapsto \mathtt{k}_\mathtt{k})$ are in S or both $\neg(\mathtt{n}_\mathtt{k} \longmapsto \mathtt{k}_\mathtt{k})$ and $\mathtt{k}\text{'s } \mathtt{n}_\mathtt{k} \longmapsto \mathtt{k}_\mathtt{k}$ are in S. This means that S is inconsistent with the axiom Current Principal. Define the local assignment l so that $l(\mathtt{k}, \mathtt{n}) = \{\mathtt{k}' : \mathtt{k}\text{'s } \mathtt{n} \longmapsto \mathtt{k}' \in S\}$. Similar to the case for $\mathrm{AX}_{\mathit{inf}}$, define the world $w = (\beta, c)$ by taking $\beta(\mathtt{g}) = \{\mathtt{k} \in K \mid \mathtt{g} \longmapsto \mathtt{k} \in S^+\}$ and defining $c(\mathtt{k})$, for each key $\mathtt{k}$, to be the set of formulas $\mathtt{n} \longmapsto \mathtt{p}$ such that $\mathtt{k}\ \mathit{cert}\ (\mathtt{n} \longmapsto \mathtt{p}) \in S$.

Now we have the following analogue to Lemma A.8.

Lemma A.9: *For all formulas $\psi \in \mathit{Sub}(\phi) \cup F$, we have $\psi \in S$ iff $w, l, \mathtt{k}_0 \models_o \psi$.*

Proof: Again we first show that by induction on the structure of $\psi \in \mathrm{Sub}(\phi) \cup E$ that $\psi \in S$ iff $w, l, \mathtt{k}_0 \models \psi$, and then show that the assignment l is consistent with w.

It is immediate from the construction of w that $w, l, \mathtt{k}_0 \models \psi$ iff $\psi \in S$ for ψ of the form $\mathtt{k}\ \mathit{cert}\ (\mathtt{n} \longmapsto \mathtt{p})$.

We next show that the result holds if ψ is of the form $\mathtt{p} \longmapsto \mathtt{k}'$, for $\mathtt{p} \in P$, by induction on the structure of $\mathtt{p}$. We strengthen the induction hypothesis to also show that $w, l, \mathtt{k}_0 \models \mathtt{k}\text{'s } \mathtt{p} \longmapsto \mathtt{k}'$ iff $\mathtt{k}\text{'s } \mathtt{p} \longmapsto \mathtt{k} \in S$. If $\mathtt{p}$ is a key $\mathtt{k}_1$, then $w, l, \mathtt{k}_0 \models \mathtt{k}_1 \longmapsto \mathtt{k}'$ iff $\mathtt{k}' = \mathtt{k}_1$ and by Reflexivity and Key Distinctness, $\mathtt{k}_1 \longmapsto \mathtt{k}' \in S$ iff $\mathtt{k}_1 = \mathtt{k}'$. Similarly, $w, l, \mathtt{k}_0 \models \mathtt{k}\text{'s } \mathtt{k}_1 \longmapsto \mathtt{k}'$ iff $w, l, \mathtt{k}_0 \models \mathtt{k}_1 \longmapsto \mathtt{k}'$ iff $\mathtt{k}_1 \longmapsto \mathtt{k}' \in S$ iff $\mathtt{k}\text{'s } \mathtt{k}_1 \longmapsto \mathtt{k}' \in S$, by Transitivity, Key Globality, and Converse of Globality (using the fact that S is $\mathrm{AX}_{\mathit{fin}}$-consistent).

If $\mathtt{p}$ is a global identifier $\mathtt{g}$, $w, l, \mathtt{k}_0 \models \mathtt{g} \longmapsto \mathtt{k}'$ iff $\mathtt{g} \longmapsto \mathtt{k}' \in S$ by the definition of β. The argument for $\mathtt{k}\text{'s } \mathtt{g} \longmapsto \mathtt{k}'$ is identical to the case that $\mathtt{p} = \mathtt{k}$.

If $\mathtt{p}$ is the local name $\mathtt{n}$, then $w, l, \mathtt{k}_0 \models \mathtt{n} \longmapsto \mathtt{k}'$ iff $\mathtt{k}' \in l(\mathtt{k}_0, \mathtt{n})$ iff $\mathtt{k}_0\text{'s } \mathtt{n} \longmapsto \mathtt{k}' \in S$ iff $\mathtt{n} \longmapsto \mathtt{k}' \in S$, by choice of $\mathtt{k}_0$. Similarly, $w, l, \mathtt{k}_0 \models \mathtt{k}\text{'s } \mathtt{n} \longmapsto \mathtt{k}'$ iff $\mathtt{k}' \in l(\mathtt{n}, \mathtt{k})$ iff $\mathtt{k}\text{'s } \mathtt{n} \longmapsto \mathtt{k}' \in S$.

Finally, if $\mathtt{p}$ is of the form $\mathtt{q}\text{'s } \mathtt{r}$, then $w, l, \mathtt{k}_0 \models \mathtt{q}\text{'s } \mathtt{r} \longmapsto \mathtt{k}'$ iff there exists a key $\mathtt{k}''$ such that $w, l, \mathtt{k}_0 \models \mathtt{q} \longmapsto \mathtt{k}''$ and $w, l, \mathtt{k}_0 \models (\mathtt{k}'')\text{'s } \mathtt{r} \longmapsto \mathtt{k}'$ iff (by the induction hypothesis) there exists a key $\mathtt{k}''$ such that $\mathtt{q} \longmapsto \mathtt{k}'' \in S$ and $(\mathtt{k}'')\text{'s } \mathtt{r} \longmapsto \mathtt{k}' \in S$ iff $\mathtt{q}\text{'s } \mathtt{r} \longmapsto \mathtt{k}' \in S$. The "only if" direction of the last equivalence follows using Left Monotonocity and Transitivity; the "if" direction follows from Witnesses. The argument for $\mathtt{k}$'s ($\mathtt{q}$'s $\mathtt{r}$) $\longmapsto \mathtt{k}'$ is identical, using Associativity: $w, l, \mathtt{k}_0 \models \mathtt{k}$'s ($\mathtt{q}$'s $\mathtt{r}$) $\longmapsto \mathtt{k}'$ iff there exists a key $\mathtt{k}''$ such that $w, l, \mathtt{k}_0 \models \mathtt{k}$'s $\mathtt{q} \longmapsto \mathtt{k}''$ and $w, l, \mathtt{k}_0 \models$ ($\mathtt{k}''$)'s $\mathtt{r} \longmapsto \mathtt{k}'$ iff there exists a key $\mathtt{k}''$ such that $\mathtt{k}$'s $\mathtt{q} \longmapsto \mathtt{k}'' \in S$ and ($\mathtt{k}''$)'s $\mathtt{r} \longmapsto \mathtt{k}' \in S$ iff $\mathtt{k}$'s ($\mathtt{q}$'s $\mathtt{r}$) $\longmapsto \mathtt{k}' \in S$.

We now continue with our induction in the case that $\mathtt{p} \longmapsto \mathtt{q}$. Note that $w, l, \mathtt{k}_0 \models \mathtt{p} \longmapsto \mathtt{q}$ iff $w, l, \mathtt{k}_0 \models \mathtt{q} \longmapsto \mathtt{k}'$ implies $w, l, \mathtt{k}_0 \models \mathtt{p} \longmapsto \mathtt{k}'$ for all $\mathtt{k}' \in K$ iff (by the induction hypothesis) $\mathtt{q} \longmapsto \mathtt{k}' \in S$ implies $\mathtt{p} \longmapsto \mathtt{k}' \in S$ iff $\mathtt{p} \longmapsto \mathtt{q} \in S$. The "only if" direction of the last equivalence follows immediately from Transitivity; the "if" direction follows from Witnesses.

We complete the induction proof by observing that if ψ is of the form $\neg\psi$ or $\psi_1 \wedge \psi_2$, the result follows immediately from the induction hypothesis.

To show that l is consistent with w, suppose that $\mathtt{n} \longmapsto \mathtt{p} \in c(\mathtt{k})$. By construction, this means that $\mathtt{k}$ *cert* $(\mathtt{n} \longmapsto \mathtt{p}) \in S$. By Key Linking, we must also have $\mathtt{k}$'s $\mathtt{n} \longmapsto \mathtt{k}$'s $\mathtt{p} \in S$. By what we have just shown, $w, l, \mathtt{k}_0 \models \mathtt{k}$'s $\mathtt{n} \longmapsto \mathtt{k}$'s $\mathtt{p}$. It follows that $w, l, \mathtt{k} \models \mathtt{n} \longmapsto \mathtt{p}$. Thus, l is consistent with w. ∎

This completes the proof of Theorem 3.2 in the case that K is finite. Note that since we can assume without loss of generality that $|K| \leq 2 \cdot |\phi|^2$ here (otherwise the argument for the case that K is infinite applies) the proof also shows that Proposition 3.3 holds. ∎

Theorem 3.5: *The same formulas are c-valid and o-valid; i.e., for all formulas ϕ, we have $\models_{\mathrm{o}} \phi$ iff $\models_{\mathrm{c}} \phi$.*

Proof: We show that $\neg\phi$ is o-satisfiable iff $\neg\phi$ is c-satisfiable, which is equivalent to the claim. The direction from c-satisfiability to o-satisfiability is straightforward: Since for every world w the local name assignment l_w is w-consistent, it follows from $w, \mathtt{k} \models_{\mathrm{c}} \neg\phi$ that $w, l_w, \mathtt{k} \models_{\mathrm{o}} \neg\phi$. Thus, it remains to show that if $\neg\phi$ is o-satisfiable, then it is c-satisfiable.

So suppose that $\neg\phi$ is o-satisfiable. By Proposition 3.3, there is a world $w = (\beta, c)$, local name assignment l, and principal $\mathtt{k}$ such that $w, l, \mathtt{k} \models_o \neg\phi$ and a finite subset K' of K such that $l(\mathtt{k}', \mathtt{n}) \subseteq K'$ for all $\mathtt{k}' \in K$ and $\mathtt{n} \in N$, and $\beta(\mathtt{g}) \subseteq K'$ for all global names $\mathtt{g}$. By standard propositional reasoning, $\neg\phi$ is equivalent to a disjunctive normal form expression in which the atoms are of the form $\mathtt{p} \longmapsto \mathtt{q}$ and $\mathtt{k}_1$ *cert* ψ, where $\mathtt{p}$ and $\mathtt{q}$ are principal expressions, $\mathtt{k}_1$ is a key, and ψ is a formula. If $w, l, \mathtt{k} \models_{\mathrm{o}} \neg\phi$ then one of the disjuncts σ is satisfied, i.e., $w, l, \mathtt{k} \models_{\mathrm{o}} \sigma$. Suppose that σ is the conjunction of the formulas in the set $A \cup B$, where

1. A is a set of formulas of the form $\mathtt{p} \longmapsto \mathtt{q}$ or $\neg(\mathtt{p} \longmapsto \mathtt{q})$,
2. B is a set of formulas of the form $\mathtt{k}_1$ *cert* ψ or $\neg(\mathtt{k}_1$ *cert* $\psi)$.

Let K_ϕ be the set of keys that appear in the formula ϕ together with K' and $\mathtt{k}$. Let N_ϕ be the set of local names that appear in ϕ. Define the world $w' = (\beta', c')$ as follows. Take the interpretation of global names β' to be equal to β, the interpretation of global names in w. Define c' by taking the set of certificates $c'(\mathtt{k}')$ to be the empty if $\mathtt{k}' \notin K_\phi$ and to consist of $c(\mathtt{k}')$ together with all certificates of the form $\mathtt{n} \longmapsto \mathtt{p}_{\mathtt{k}'',\phi}$ if $\mathtt{k}' \in K_\phi$, $\mathtt{n} \in N_\phi$, and $\mathtt{k}'' \in l(\mathtt{n}, \mathtt{k}')$, where $\mathtt{p}_{\mathtt{k}'',\phi}$ is a principal expression of the form ($\mathtt{k}''$)'s ($\mathtt{k}''$)'s $\ldots$ ($\mathtt{k}''$) that does not appear in

ϕ. (Clearly we can make the expression sufficiently long so as to ensure it does not appear in ϕ.) Clearly $\cup_{\mathtt{k}' \in K} c(\mathtt{k}')$ is finite.

We show that $w', \mathtt{k} \models_{\mathrm{c}} \sigma$. It follows from this that $w', \mathtt{k} \models_{\mathrm{c}} \neg\phi$. Note first that from the fact that $c(\mathtt{k}') \subseteq c'(\mathtt{k}')$ for all $\mathtt{k}'$, it follows that $w', \mathtt{k} \models_{\mathrm{c}} \mathtt{k}'$ *cert* ψ for all formulas $\mathtt{k}'$ *cert* ψ in B. Moreover, if $\neg(\mathtt{k}'$ *cert* $\psi)$ is in B then, since the expressions $\mathtt{p}_{\mathtt{k}'',\phi}$ on the right-hand side of the certificates in $c'(\mathtt{k}') - c(\mathtt{k})$ do not appear in ϕ it follows that $w', \mathtt{k} \models_{\mathrm{c}} \neg(\mathtt{k}'$ *cert* $\psi)$. Thus $w', \mathtt{k} \models_{\mathrm{c}} B$.

It remains to show that the formulas in A are satisfied. To show this, we show that

$$l_{w'}(\mathtt{n}, \mathtt{k}') = l(\mathtt{n}, \mathtt{k}') \text{ for all } \mathtt{n} \in N_\phi \text{ and } \mathtt{k}' \in K_\phi. \tag{2}$$

It easily follows from (2), the fact that all keys in ϕ are in K', and the fact that global names have the same interpretation in w and w' that $[\![p]\!]_{w', l_{w'}, \mathtt{k}'} = [\![p]\!]_{w, l, \mathtt{k}'}$ for all principal expressions p occurring in A and all keys $\mathtt{k}' \in K_\phi$. This in turn is easily seen to imply that $w', \mathtt{k} \models_{\mathrm{c}} A$.

It remains to prove (2). It is almost immediate from the definition of l' that $l_{w'}(\mathtt{n}, \mathtt{k}') \supseteq l(\mathtt{n}, \mathtt{k}')$ for all $\mathtt{n} \in N_\phi$ and $\mathtt{k}' \in K_\phi$. For the opposite containment, we prove by induction on j that $(T_{w'} \uparrow j)(\mathtt{n}, \mathtt{k}') \subseteq l(\mathtt{n}, \mathtt{k}')$ for all $j \in \mathbf{N}$, $\mathtt{n} \in N_\phi$, and $\mathtt{k}' \in K_\phi$. The base case $j = 0$ is trivial. For the induction step, suppose that $j = j' + 1$ and $\mathtt{k}'' \in (T_{w'} \uparrow j)(\mathtt{n}, \mathtt{k}')$. Thus, $\mathtt{k}'' \in (T_{w'}(T_{w'} \uparrow j'))(\mathtt{n}, \mathtt{k}')$, which means that $\mathtt{k}'' \in [\![\mathtt{p}]\!]_{w', T_{w'} \uparrow j', \mathtt{k}'}$ for some principal expression $\mathtt{p}$ such that $\mathtt{n} \longmapsto \mathtt{p} \in c'(\mathtt{k}')$. There are two possibilities: (1) $\mathtt{n} \longmapsto \mathtt{p} \in c(\mathtt{k}')$ or (2) $\mathtt{n} \longmapsto \mathtt{p} \in c'(\mathtt{k}') - c(\mathtt{k}')$. In case (2), $\mathtt{p}$ must be of the form $\mathtt{p}_{\mathtt{k}_1,\phi}$ so $[\![\mathtt{p}]\!]_{w', T_{w'} \uparrow j', \mathtt{k}'} = \{\mathtt{k}_1\}$ and $\mathtt{k}_1 = \mathtt{k}''$. But in this case, by construction, $\mathtt{k}'' \in l(\mathtt{n}, \mathtt{k}')$. In case (1), using the induction hypothesis and the fact that global names and keys in $\mathtt{p}$ have the same interpretation in w and w' (this interpretation being a subset of K'), we get that $[\![\mathtt{p}]\!]_{w', T_{w'} \uparrow j', \mathtt{k}'} \subseteq [\![\mathtt{p}]\!]_{w, l, \mathtt{k}'}$. Thus, $\mathtt{k}'' \in [\![\mathtt{p}]\!]_{w, l, \mathtt{k}'}$. Because l is w-consistent and $\mathtt{n} \longmapsto \mathtt{p} \in c(\mathtt{k}')$, we again obtain that $\mathtt{k}'' \in l(\mathtt{n}, \mathtt{k}')$, as required.

Since $l_{w'}(\mathtt{n}, \mathtt{k}')$ is the union of the $(T_{w'} \uparrow j)(\mathtt{n}, \mathtt{k}')$, it follows that $l_{w'}(\mathtt{n}, \mathtt{k}') = l(\mathtt{n}, \mathtt{k}')$. This completes the proof of (2). ∎

Proposition 3.8: *Let Γ be any c-satisfiable boolean combination of formulas of the form* $\mathtt{k}$ *cert* ϕ, *and let Δ be any boolean combination of formulas of the form* $\mathtt{p} \longmapsto \mathtt{q}$ *where neither* $\mathtt{p}$ *nor* $\mathtt{q}$ *contains a local name. Then $\models_{\mathrm{c}} \Gamma \Rightarrow \Delta$ iff $\models_{\mathrm{c}} \Delta$.*

Proof: Clearly $\models_{\mathrm{c}} \Delta$ implies $\models_{\mathrm{c}} \Gamma \Rightarrow \Delta$. For the converse, suppose by way of contradiction that $\models_{\mathrm{c}} \Gamma \Rightarrow \Delta$ and there is a world $w = (\beta, c)$ and a principal $\mathtt{k}$ such that $w, \mathtt{k} \models_{\mathrm{c}} \neg\Delta$. Since Γ is assumed to be c-satisfiable, there exists a world $w' = (\beta', c')$ and a principal $\mathtt{k}'$ such that $w', \mathtt{k}' \models_{\mathrm{c}} \Gamma$. Let w'' be the world (β, c'). Then a straightforward induction shows that for all principal expressions $\mathtt{p}$ not containing a local name, we have $[\![\mathtt{p}]\!]_{w'', l_{w''}, \mathtt{k}} = [\![\mathtt{p}]\!]_{w, l_w, \mathtt{k}}$. Moreover, for all keys $\mathtt{k}_1$ and formulas ϕ, we have $w'', \mathtt{k} \models_{\mathrm{c}} \mathtt{k}_1$ *cert* ϕ iff $w', \mathtt{k}' \models_{\mathrm{c}} \mathtt{k}_1$ *cert* ϕ. It follows that $w'', \mathtt{k} \models_{\mathrm{c}} \Gamma \wedge \neg\Delta$, giving us our desired contradiction. ∎

Theorem 4.1: *Suppose $\mathtt{k}_1, \mathtt{k}_2$ are principals, $w = (\beta, c)$ is a world, and* $\mathtt{p}$ *is a principal expression. Let E_w be the set of all the formulas* $\mathtt{g} \longmapsto \mathtt{k}$ *for all global names* $\mathtt{g}$ *and keys* $\mathtt{k} \in \beta(\mathtt{g})$ *and the formulas* $\mathtt{k}$ *cert* ϕ *for all keys* $\mathtt{k}$ *and formulas* $\phi \in c(\mathtt{k})$. *The following are equivalent:*

1. $\mathtt{k}_1 \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{p})$,

2. $w, \mathtt{k}_2 \models_{\mathrm{c}} \mathtt{p} \longmapsto \mathtt{k}_1$,

3. $w', \mathtt{k}_2 \models_{\mathrm{c}} \mathtt{p} \longmapsto \mathtt{k}_1$ *for all worlds* $w' \geq w$,

4. $E_w \models_c \mathtt{k}_2$*'s* $\mathtt{p} \longmapsto \mathtt{k}_1$,

5. $E_w \models_o \mathtt{k}_2$*'s* $\mathtt{p} \longmapsto \mathtt{k}_1$.

Proof: The presentation of REF2 in Figure 1 is still slightly informal, combining recursion and nondeterminism. To make it fully precise, define a *computation tree* of REF2 to be a finite tree labelled by expressions of the form "$\mathtt{k}_1 \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{p})$", such that if N is a node so labelled, then one of the following four conditions holds:

1. $\mathtt{p}$ is a key $\mathtt{k}$, we have $\mathtt{k} = \mathtt{k}_1 = \mathtt{k}_2$, and N is a leaf of the tree,
2. $\mathtt{p}$ is a global name $\mathtt{g}$ and $\mathtt{k}_1 \in \beta(\mathtt{g})$,
3. $\mathtt{p}$ is a local name $\mathtt{n}$ and $c(\mathtt{k}_2)$ contains a formula $\mathtt{n} \longmapsto \mathtt{q}$ and N has exactly one child, labelled "$\mathtt{k}_1 \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{q})$",
4. $\mathtt{p}$ is of the form $\mathtt{q}$'s $\mathtt{r}$ and N has exactly two children, labelled "$\mathtt{k} \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{q})$" and "$\mathtt{k}_1 \in \mathrm{REF2}(\mathtt{k}, \beta, c, \mathtt{r})$", for some key $\mathtt{k}$.

We take $\mathtt{k}_1 \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{p})$ to mean that there exists a computation tree of REF2 with root labelled "$k_1 \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{p})$".

Given a world $w = (\beta, c)$ and $m \in \mathbf{N}$, let $l_m = T_w \uparrow m$. The following result establishes a correspondence between the stages of the computation of l_w and the computation trees of REF2. The proof is by a straightforward induction on m, with a subinduction on the structure of $\mathtt{p}$.

Lemma A.10: *For all* $m \in \mathbf{N}$, *keys* $\mathtt{k}_1, \mathtt{k}_2$, *worlds* $w = (\beta, c)$, *and principal expressions* $\mathtt{p}$, *we have* $\mathtt{k}_1 \in [\![\mathtt{p}]\!]_{w, l_m, \mathtt{k}_2}$ *iff there exists a computation tree of* REF2 *of height at most* m *whose root is labelled* "$\mathtt{k}_1 \in \mathrm{REF2}(\mathtt{k}_2, \beta, c, \mathtt{p})$".

Using the fact that $l_w = \sqcup\{l_m : m \in \mathbf{N}\}$, Lemma A.1, and Lemma A.10, we obtain the equivalence between (1) and (2).

The proof of the implication from (2) to (3) is by a straightforward induction on the structure of $\mathtt{p}$; that is, for fixed $w' \geq w$, we show by induction on the structure of $\mathtt{p}$ that if $w, \mathtt{k}_2 \models_{\mathrm{c}} \mathtt{p} \longmapsto \mathtt{k}_1$ then $w', \mathtt{k}_2 \models_{\mathrm{c}} \mathtt{p} \longmapsto \mathtt{k}_1$. The opposite implication from (3) to (2) is trivial, since $w \geq w$. For the implication from (3) to (4), suppose that (3) holds and (4) does not. Then for some world w' and key $\mathtt{k}$ we have $w', \mathtt{k} \models_{\mathrm{c}} E_w$ and $w', \mathtt{k} \models_{\mathrm{c}} \neg(\mathtt{k}_2\text{'s } \mathtt{p} \longmapsto \mathtt{k}_1)$. The latter implies $w', \mathtt{k}_2 \models_{\mathrm{c}} \neg(\mathtt{p} \longmapsto \mathtt{k}_1)$. Since $w', \mathtt{k} \models_{\mathrm{c}} E_w$, it follows that $w' \geq w$. Thus, by (3), $w', \mathtt{k}_2 \models_{\mathrm{c}} \mathtt{p} \longmapsto \mathtt{k}_1$, contradicting our assumption. The implication from (4) to (3) is immediate, since $w', \mathtt{k}_2 \models_{\mathrm{c}} E_w$ for all $w' \geq w$. Finally, the equivalence between (4) and (5) is just a special case of Theorem 3.5. ∎

Proposition 5.1: *If M represents w and l then for all principal expressions* $\texttt{p}$ *and* $x, y \in K \cup G \cup N$ *we have* $M \models \tau_{x,y}(\texttt{p})$ *iff* $x, y \in K$ *and* $w, l, x \models \texttt{p} \longmapsto y$.

Proof: By a straightforward induction on the structure of $\texttt{p}$. The base cases, where $\texttt{p} \in K \cup G \cup N$, are immediate from the definition of "represents" and the semantics of the logic. The inductive case, where $\texttt{p} = \texttt{q}$'s $\texttt{r}$, is immediate from the semantics and the definition of the translation. ∎

Theorem 5.2: *The minimal Herbrand model M_w of Σ_w represents w and l_w.*

Proof: (Sketch) The proof proceeds by showing a direct correspondence between the construction of the minimal Herbrand model of Σ_w and the fixpoint construction of l_w.

The theory of logic programming [Llo87] associates with the Horn theory Σ_w an operator Φ_w on the space of Herbrand models on the vocabulary V, defined by $\texttt{name}(x, y, z) \in \Phi_w(M)$ if there exists a substitution instance of a formula in Σ_w of the form $B \Rightarrow \texttt{name}(x, y, z)$ such that $M \models B$. The least Herbrand model M_w of Σ_w is then equal to $\Phi_w \uparrow \omega = \bigcup_{m \in \mathbf{N}} \Phi_w \uparrow m$, where $\Phi_w \uparrow 0 = \emptyset$ and $\Phi_w \uparrow m + 1 = \Phi_w(\Phi_w \uparrow m)$ for $m \geq 0$.

Let T_w be the operator on local name assignments defined in the proof of Theorem 3.1. Using Proposition 5.1 to handle the rules in Σ_w corresponding to certificates, we may then show by a straightforward induction on m that for all $m \geq 1$, the Herbrand model $\Phi \uparrow m$ represents the world w and the local name assignment $T_w \uparrow m$. It follows that $M_w = \Phi \uparrow \omega$ represents $l_w = T_w \uparrow \omega$. ∎

Theorem 6.1: *AX^{self}_{inf} (resp., AX^{self}_{fin}) is a sound and complete axiomatization of LLNC^s with respect to the open semantics if K is infinite (resp., K is finite).*

Proof: The argument is very similar to that in the proof of Theorem 3.2. First suppose that K is infinite.

We add the following clauses to the definition of P:

6. $\texttt{Self} \in P$,

7. if $\texttt{n} \in P$ is a local name then $\texttt{Self}$'s $\texttt{n} \in P$.

We also add the following clauses to the definition of S^+, corresponding to the new axioms for $\texttt{Self}$.

(ClSP) if $\texttt{Self}$'s $\texttt{p} \in P$ then $\texttt{Self}$'s $\texttt{p} \longmapsto \texttt{p} \in S^+$ and $\texttt{p} \longmapsto \texttt{Self}$'s $\texttt{p} \in S^+$,

(ClPS) if $\texttt{p}$'s $\texttt{Self} \in P$ then $\texttt{p}$'s $\texttt{Self} \longmapsto \texttt{p} \in S^+$ and $\texttt{p} \longmapsto \texttt{p}$'s $\texttt{Self} \in S^+$,

(ClSE) if $\texttt{Self} \longmapsto \texttt{p} \in S^+$ and $\texttt{p}$'s $\texttt{k} \longmapsto \texttt{k} \in S^+$ then $\texttt{p} \longmapsto \texttt{Self} \in S^+$.

Lemma A.5 still applies. The definitions following this lemma, up to and including that of S^* are unchanged. However, the construction of the model changes slightly. We no longer use $\texttt{k}_0$ to represent the "current principal", instead, we use the key $\texttt{k}_*$ that the construction associates with $\texttt{Self}$. This could be either a key in P_1 or one of the keys $\texttt{k}_c$ for $c \in O$, depending on whether $\texttt{Self}$ is key-equivalent or open. Note that we cannot have $\texttt{Self}$ empty (thanks to the Identity axiom). If $\texttt{Self}$ is key-equivalent, then by (ClKD) it is equivalent to at most one key $\texttt{k} \in P$. In this case, we define $\texttt{k}_* = \texttt{k}$. If $\texttt{Self}$ is open we define $\texttt{k}_*$ to be $\texttt{k}_c$, where $c = [\texttt{Self}]$.

We now define w and l exactly as before, except that we now set $l(\texttt{k}_0, \texttt{n}) = \emptyset$, since we no longer use $\texttt{k}_0$ as the "current principal." The following lemma is the analogue of Lemma A.7.

Lemma A.11: *For all expressions* $\texttt{p} \in P$*, we have* $[\![\texttt{p}]\!]_{w,l,\texttt{k}_*} = I(\texttt{p})$*.*

Proof: The proof is very similar to that of Lemma A.7; we just describe the modifications required. The base cases for $\texttt{p}$ a global name or a key are identical.

When $\texttt{p} = \texttt{n}$ is a local name, we proceed as follows. There are two possibilities, depending on whether $\texttt{k}_* \in P$ or not. Suppose first that $\texttt{k}_* \in P$. Then we have $\texttt{k}_* \approx \texttt{Self}$ and, by (ClLM) and (ClSP), $\texttt{k}_*$'s $\texttt{n} \approx \texttt{Self}$'s $\texttt{n} \approx \texttt{n}$. It then follows by (ClT) and construction of l that $\texttt{n} \longmapsto \texttt{k} \in S^*$ iff $\texttt{k}_*$'s $\texttt{n} \longmapsto \texttt{k} \in S^*$ iff $\texttt{k} \in l(\texttt{k}_*, \texttt{n})$, as required.

If $\texttt{k}_* = \texttt{k}_c$ for c an open class, we proceed as follows. If $\texttt{k} \in I(\texttt{n})$, then we consider two cases, depending on whether $\texttt{k} \in P_1$. If $\texttt{k} \in P_1$, then $\texttt{n} \longmapsto \texttt{k} \in S^+$ and it follows that $\texttt{Self}$'s $\texttt{n} \longmapsto \texttt{k}$ by (ClSP) and (ClT). Since $\texttt{Self} \approx \texttt{Self}$ it is immediate that $\texttt{k} \in [\![\texttt{p}]\!]_{w,l,\texttt{k}_*}$. Alternatively, if $\texttt{k} = \texttt{k}_d$, for $d \in O$, then we have $\texttt{n} \longmapsto \texttt{q} \in S^+$ for some $\texttt{q} \in d$. By (ClSP) and (ClT) it follows that $\texttt{Self}$'s $\texttt{n} \longmapsto \texttt{q} \in S^+$, hence $\texttt{Self}$'s $\texttt{n} \longmapsto \texttt{k} \in S^*$. As before, this implies that $\texttt{k} \in [\![\texttt{n}]\!]_{w,l,\texttt{k}_*}$.

For the opposite inclusion, suppose that $\texttt{k} \in [\![\texttt{n}]\!]_{w,l,\texttt{k}_*}$. Since we are assuming that $\texttt{Self}$ is open, there must be some $\texttt{q} \approx \texttt{Self}$ such that $\texttt{q}$'s $\texttt{n} \longmapsto \texttt{k} \in S^*$. By (ClLM), we have $\texttt{Self}$'s $\texttt{n} \longmapsto \texttt{q}$'s $\texttt{n} \in S^+$. It follows using (ClT) that $\texttt{Self}$'s $\texttt{n} \longmapsto \texttt{k} \in S^*$, hence $\texttt{n} \longmapsto \texttt{k} \in S^*$. This completes the argument for the base case of $\texttt{n}$ a local name.

There is now an additional base case for $\texttt{p} = \texttt{Self}$. Here, note that $[\![\texttt{Self}]\!]_{w,l,\texttt{k}_*} = \{\texttt{k}_*\}$. We therefore need to show that $\texttt{Self} \longmapsto \texttt{k} \in S^*$ iff $\texttt{k} = \texttt{k}_*$. When $\texttt{k}_* \in P_1$, we have $\texttt{Self} \approx \texttt{k}_*$, so $\texttt{Self} \longmapsto \texttt{k} \in S^*$ iff $\texttt{k}_* \longmapsto \texttt{k}$, and the claim follows by (ClKD) and (ClT) as in the base case for keys. The alternative is that $\texttt{k}_* = \texttt{k}_c$ for $c = [\texttt{Self}] \in O$. Since have $\texttt{Self} \longmapsto \texttt{k}_c \in S^*$ by construction of S^*, it remains to prove that if $\texttt{Self} \longmapsto \texttt{k} \in S^*$ then $\texttt{k} = \texttt{k}_c$. Now we cannot have $\texttt{Self} \longmapsto \texttt{k} \in S^*$ for $\texttt{k} \in P_1$, for then by the argument above that $\texttt{Self}$ is nonempty and (ClSE), we have $\texttt{k} \longmapsto \texttt{Self} \in S^+$, contradicting the assumption that c is open. Thus, we must have $\texttt{k} = \texttt{k}_d$ for some $d \in O$. In this case, there exists $\texttt{q} \in d$ such that $\texttt{Self} \longmapsto \texttt{q} \in S^+$. Since d is open, we have $\texttt{q}$'s $\texttt{k}_0 \longmapsto \texttt{k}_0 \in S^+$, hence $\texttt{q} \longmapsto \texttt{Self} \in S^+$ by (ClSE). Thus, $\texttt{Self} \approx \texttt{q}$, and it follows that $d = c$, hence $\texttt{k} = \texttt{k}_*$ as required. This completes the argument for the base case where $\texttt{p} = \texttt{Self}$.

The inductive case is exactly as before, except that we need to consider the new case $\texttt{p}$'s $\texttt{Self}$. Here, we note that $[\![\texttt{p}\text{'s }\texttt{Self}]\!]_{w,l,\texttt{k}_*} = [\![\texttt{p}]\!]_{w,l,\texttt{k}_*}$. Thus, by the induction hypothesis, we are required to prove that $\texttt{p} \longmapsto \texttt{k} \in S^*$ iff $\texttt{p}$'s $\texttt{Self} \longmapsto \texttt{k} \in S^*$. This follows using (ClPS) and (ClT). ∎

The remainder of the proof in the case that K is infinite proceeds as before, using $\texttt{k}_*$ in place of $\texttt{k}_0$.

If K is finite, the proof is even closer to that for the logic without $\texttt{Self}$. As sketched in the main text, because S is consistent, it follows from Identity, Witnesses, and Self-is-key that there must be some key $\mathtt{k}_* \in K$ such that $\mathtt{Self} \longmapsto \mathtt{k}_* \in S$. For this key $\mathtt{k}_*$, we must have $\mathtt{k}_*$'s $\mathtt{n} \longmapsto \mathtt{k} \in S$ iff $\mathtt{n} \longmapsto \mathtt{k} \in S$. Thus, $\mathtt{k}_*$ plays the role of $\mathtt{k}_0$ in the earlier argument. (Note that we now no longer need Current Principal to ensure the existence of $\mathtt{k}_0$.) The rest of the argument is unchanged.) ∎

Acknowledgments

Work on this paper was done while the second author was with the School of Computing Sciences, University of Technology, Sydney. This work was supported in part by NSF under grant IRI-96-25901 and by a UTS internal research grant. A preliminary version of this paper appeared in the *Proceedings of the 12th IEEE Computer Security Foundations Workshop*, 1999, pp. 111–122.

References

[Aba98] M. Abadi. On SDSI's linked local name spaces. *Journal of Computer Security*, 6(1-2):3–21, 1998.

[ABLP93] M. Abadi, M. Burrows, B. Lampson, and G. D. Plotkin. A calculus for access control in distributed systems. *ACM Transactions on Programming Languages and Systems*, 15(4):706–734, 1993.

[BFL96] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proceedings 1996 IEEE Symposium on Security and Privacy*, pages 164–173, 1996.

[Bir67] G. Birkhoff. *Lattice Theory.* American Mathematical Society, Providence, R.I., 3rd edition edition, 1967.

[EK76] M.H. van Emden and R. A. Kowalski. The semantics of predicate logic as a programming language. *Journal of the ACM*, 23(4):733–742, 1976.

[FHMV95] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* MIT Press, Cambridge, Mass., 1995.

[GH93] A. J. Grove and J. Y. Halpern. Naming and identity in propositional logics, Part I: the propositional case. *Journal of Logic and Computation*, 3(4):345–378, 1993.

[Gro98] SPKI Working Group. Simple public key infrastructure, internet draft. at `http://www.ietf.org/html.charters/spki-charter.html`, 1998.

[HM90] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, 1990. A preliminary version appeared in *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, 1984.

[HvdM99] J. Y. Halpern and R. van der Meyden. Adding revocation and timestamps to a logic for sdsi's linked local name spaces. unpublished manuscript, 1999.

[HvdMS99] J.Y. Halpern, R. van der Meyden, and F. Schneider. Logical foundations for trust management. manuscript, 1999.

[LABW92] B. Lampson, M. Abadi, M. Burrows, and E. Wobber. Authentication in distributed systems: Theory and practice. *ACM Transactions on Computer Systems*, 10(4):265–310, 1992.

[Llo87] J. W. Lloyd. *Foundations of Logic Programming.* Springer-Verlag, Berlin, 2nd edition edition, 1987.

[LNS82] J.-L. Lassez, V. L. Nguyen, and E. A. Sonenberg. Fixed point theorems and semantics: a folk tale. *Information Processing Letters*, 14(3):112–116, 1982.

[RL96] R.L. Rivest and B. Lampson. SDSI — a simple distributed security infrastructure. at `http://theory.lcs.mit.edu/∼cis/sdsi.html`, 1996.

[Ull88] J. D. Ullman. *Principles of Database and Knowledge Base Systems, Volume I.* Computer Science Press, 1988.

[Ull89] J. D. Ullman. *Principles of Database and Knowledge Base Systems, Volume II: The New Technologies.* Computer Science Press, 1989.