



A Sequent Calculus for Counterfactual Reasoning

McKenna McCall
Carnegie Mellon University
USA

Lay Kuan Loh
Carnegie Mellon University
USA

Limin Jia
Carnegie Mellon University
USA

ABSTRACT

Counterfactual conditions such as “if A were not true, then C would not have been true” have been formally studied by philosophers for causal claims for decades. Counterfactuals are often used informally in practice for diagnosing systems and identifying errors or misconfigurations. This paper develops a proof theory for counterfactual reasoning of Horn clauses, which have applications in domains including security and database and program analysis. The application to security that this paper focuses on is modeling and reasoning about probing attacks in Datalog-based trust management systems, where an attacker can apply counterfactual reasoning to obtain sensitive information embedded in the system.

Our work is inspired by a Hilbert-style axiomatized system for counterfactual reasoning for Horn clauses, which are hard to use to construct proofs or study properties of the system. To alleviate this difficulty, we develop a sequent calculus from first principles. We show that the sequent calculus has cut elimination and is sound and complete with regard to the corresponding Hilbert style axiomatized system. We also show how to construct proofs that model practical counterfactual reasoning scenarios in trust management systems using our sequent calculus rules.

KEYWORDS

counterfactual reasoning, access control, Horn clauses, sequent calculus, cut elimination

1 INTRODUCTION

Counterfactual conditions such as “if A were true, then C would have been true” and “if A were not true, then C would not have been true” describe what could or would happen under different conditions. For decades, philosophers have investigated using counterfactuals to explain causal claims [10, 22, 26]. In the above example, A is a likely cause of C . Counterfactual reasoning is a meta-level reasoning about a system. In practice, people often apply ad-hoc counterfactual reasoning to diagnosing and analyzing concrete systems, both for legitimate purposes, such as debugging, as well as malicious ones, such as probing secrets embedded in the system. The analyst crafts a set of inputs, feeds each input to the system, and then draws conclusions based on differences in observable outputs. Counterfactuals such as “if the input weren’t a null string,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PLAS'17, October 30, 2017, Dallas, TX, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5099-0/17/10...\$15.00

<https://doi.org/10.1145/3139337.3139342>

the system would not have crashed” are used to pinpoint possible causes of problems.

One important research question is how to formalize counterfactual reasoning to provide a solid foundation for correct uses of such reasoning in practice. In particular, we are interested in formalizing counterfactual reasoning in systems that can be modeled using Horn clauses. One concrete example of such systems are Datalog-based trust management systems [4, 12]. These systems typically use a set of Horn clauses to specify the attributes of users and the policies for who can access what system resources. For instance, suppose University U has a policy stating that a student can access the gym as long as the student signs up for the gym service. We can encode this policy as the following Horn clause:

$$U.\text{canAccessGym}(x) :- x.\text{registerForGym}.$$

Here, $:-$ is read as a reversed implication. We write $K.\varphi$ to mean that principal K asserts that φ is true. U is the principal representing the authority of the university. Note that this notation can be encoded as a predicate (Section 2).

One type of counterfactual reasoning in such systems is a *probing attack*. A probing attack against a trust management system involves someone who has the ability to submit additional clauses to the system. The attacker probes the system by adding clauses to the system and observing changes in access rights. The added clauses often make use of sensitive information of the system that the attacker should not have access to.

Continuing the example, an attacker A wants to know whether her friend B has a low GPA, information which is private and should not be leaked to A . A knows that she does not have access to the gym currently. A then adds the following clause into the system:

$$A.\text{registerForGym} :- U.\text{lowGPA}(B).$$

This clause states that if the University authority says that student B has a low GPA, then A registers for the gym. Given this additional clause, A gains access to the gym. From these two cases, A can deduce that it must be the case that B has a low GPA. The attacker is sure of the secret value based on counterfactual reasoning: if the secret value were different, then a particular access right would not have been granted.

This paper aims to design a formal logical system for counterfactual reasoning of Horn clauses from first principles. Existing work on counterfactual reasoning of Horn clauses provides axiomatized formulations. The advantage of a sequent calculus over a Hilbert-style axiomatized system is that it is amenable to automated theorem proving. If the sequent calculus has subformula properties, i.e., formulas in subderivations are subformulas of the conclusion, then sequent rules can be applied exhaustively to construct proofs. Axiomatized systems rely more on clever applications of axioms, so proofs can sometimes be hard to construct. Sequent calculus rules also provide an inductive principle for proving meta-properties of

the logical system. Building on prior work of a Hilbert-style axiomatized system for counterfactual reasoning for trust management policies [5], we define a sequent calculus.

We choose Horn clauses to be the subject of counterfactual reasoning because Horn clauses, due to their simplicity, have been widely used in many areas, including databases [1], logic programming [37], network verification [14, 24], program analysis [8, 38], and security [4, 7, 12]. Therefore, a solid foundation for counterfactual reasoning for Horn clauses could have impact in debugging, diagnosis, and identifying attacks in all of the above areas.

This paper makes the following technical contributions.

- We define the first sequent calculus for counterfactual reasoning of Horn clauses.
- We prove that the sequent calculus has cut elimination and that it is sound and complete with regard to an axiomatized system.
- We show how to construct proofs that demonstrate probing attacks using our sequent calculus rules.

The rest of this paper is organized as follows. We briefly review Horn clauses and trust management systems in Section 2. We present the Hilbert-style system in Section 3. Our sequent calculus rules are explained in Section 4 and properties of the calculus are discussed in Section 5. We discuss design choices and broader applications of this formalism in Section 6. Related work is presented in Section 7. Proof details may be found in our companion technical report [25].

2 BACKGROUND

A Horn clause is a first-order logical formula of the form $p_1 \wedge p_2 \cdots \wedge p_n \rightarrow q$, where each p_i and q are atomic predicates. Variables in the formula are quantified at the outermost level. We will use the Prolog/Datalog notation and write: $q :- p_1, \dots, p_n$. We call p_1 to p_n , formulas to the right of $:-$, the body of the clause, and q the head of the clause. When the domain of the variables is finite, first-order clauses can be instantiated to a finite set of propositional clauses. We write w to denote a set of propositional clauses. We write $w \Vdash_d p$ to mean that p is derivable from w , defined using the inference rules shown in Figure 1. p is derivable from a policy world if either p is in the policy world, or p is the head of a clause in the policy world, and all body propositions of that rule are derivable from the same policy world.

A trust management system can be viewed as a framework for managing distributed access control [6]. Horn clauses have been used for specifying policies in trust management systems [4, 12, 20, 23]. Below is an example encoding a Unix file system access-control policy, stating that the owner of a file can open that file and that A owns the file tmp .

$$S.\text{owns}(A, tmp). \quad S.\text{mayOpen}(x, f):- S.\text{owns}(x, f).$$

Recall from the example earlier, $S.p$ means principal S says p . The says connective was introduced in authorization logics, whose semantics and proof theory have been studied in depth (e.g., [9, 13, 15, 16]). We do not aim to model any specific authorization logic; our goal is to use the syntax to encode real-world access control examples. The says connective can be encoded by making the principal S the first argument of the predicate. Below is the

$$\frac{}{w, p \Vdash_d p} \text{At} \quad \frac{(w, p :- q_1, \dots, q_n \Vdash_d q_i)_{i \in [1, n]}}{w, p :- q_1, \dots, q_n \Vdash_d p} \text{DLOG}$$

Figure 1: Derivability relation for Datalog rules

pure Horn clause encoding.

$$\text{owns}(S, A, tmp). \quad \text{mayOpen}(S, x, f):- \text{owns}(S, x, f).$$

Role-based and attribute-based access control policies can also be encoded as a set of Horn clauses specifying the role, role assignments, attributes, and access control decisions based on roles and attributes. For the rest of this paper, we will use access control policies as our main application domain. We will use policies, rules, and clauses interchangeably. We call w a policy context.

When a principal requests access to a system resource, a proof that the principal is allowed access needs to be constructed using the clauses that encode the access control policies. From the above two policies, we can prove that A can open file tmp , using the proof system composed of the rules in Figure 1.

In trust management systems, the Horn clauses that encode the access control policies and related facts are also called *credentials*, as some systems require that they are signed using the private keys of authorities. For instance, $\text{owns}(S, A, tmp)$ is stored as a digital certificate signed by S 's private key.

Trust management systems are distributed; principals can issue their own credentials and are allowed to make their own policies. The probing attack discussed earlier is caused by the attacker issuing credentials that make use of facts that are supposed to be secret to the attacker. Based on differences in observable effects (i.e., proofs of accesses), the attacker can apply counterfactual reasoning and deduce the existence or absence of sensitive credentials.

3 COUNTERFACTUAL REASONING

We review the formalism for specifying and reasoning about counterfactuals of Horn clauses introduced by Becker et al. [5]. Our formalism uses the same syntax and semantics, but a different proof system (Section 4).

3.1 Syntax and Semantics of Counterfactual Conditions

In the example from Section 1, the probes that the attacker attempts are counterfactual conditions: if these policies were not true, access would not have been granted. Becker et al. internalized counterfactual conditions using an indexed modal operator: $\Box_w \varphi$. Informally, $\Box_w \varphi$ is true if after submitting policies in w , φ is derivable. For instance, $\Box_b :- a$ means after submitting the clause $b :- a$, b can be derived. We summarize the syntax of all the logical formulas below.

$$\begin{array}{ll} \text{Policy} & \gamma ::= \top \mid p \mid q :- \vec{p} \mid \gamma \wedge \gamma' \\ \text{Policy Context} & w ::= \top \mid w, p \mid w, q :- \vec{p} \\ \text{Formula} & \varphi ::= \gamma \mid \perp \mid \neg \varphi \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi' \\ & \quad \mid \varphi \rightarrow \varphi' \mid \Box_\gamma \varphi \end{array}$$

A policy, denoted γ , can be an atomic proposition, a Horn clause, or the conjunction of a set of Horn clauses. Formulas, denoted φ , include all of the connectives from propositional logic. The main addition is $\Box_\gamma \varphi$. It means that φ can be derived when the policy

$w \Vdash p$	iff	$w \Vdash_d p$
$w \Vdash \top$	iff	always
$w \Vdash \perp$	iff	never
$w \Vdash \neg\varphi$	iff	$w \not\Vdash \varphi$
$w \Vdash \varphi \wedge \varphi'$	iff	$w \Vdash \varphi$ and $w \Vdash \varphi'$
$w \Vdash \varphi \vee \varphi'$	iff	$w \Vdash \varphi$ or $w \Vdash \varphi'$
$w \Vdash \varphi \rightarrow \varphi'$	iff	if $w \Vdash \varphi$ then $w \Vdash \varphi'$
$w \Vdash \Box_{\gamma'} \varphi$	iff	$w, \bar{\gamma}' \Vdash \varphi$

Figure 2: Formula Semantics

γ is submitted. The Horn clauses are themselves counterfactual conditions as well: $p :- p_1, \dots, p_n$ is equivalent to $\Box_{(\wedge_{i=1}^n p_i)} p$. We will use these two notations interchangeably. A policy context w is a set of Horn clauses. The empty context is \top . We write $\wedge(w)$ to denote the policy that is the conjunction of every clause in w , and $\bar{\gamma}$ to denote the set of clauses in γ .

The example in Section 1 can be formalized as follows. φ_1 encodes A's first probe: A cannot access the gym. φ_2 is A's second probe: after submitting the policy "if B has a low GPA, then A registers for the gym", A can access the gym. γ_0 encodes the system's policy: anyone who registers for the gym can access it. γ_0 is the policy context under which the probes are carried out. As will be clear after we formally define probing attacks in Section 3.2, the attacker need not know γ_0 .

$$\begin{aligned} \varphi_1 &= \neg U.\text{canAccessGym}(A) \\ \varphi_2 &= \Box_{A.\text{registerForGym} :- U.\text{lowGPA}(B)} U.\text{canAccessGym}(A) \\ \gamma_0 &= U.\text{canAccessGym}(x) :- x.\text{registerForGym} \end{aligned}$$

A proof system (see Section 3.3) for counterfactual reasoning can show that $U.\text{lowGPA}(B)$ is the logical consequence of the conjunction of φ_1 and φ_2 .

Key definitions for the semantics of formulas are summarized in Figure 2. Formulas are interpreted under a set of policies w . The semantics for the atomic predicate is defined in terms of the derivability of Horn clauses. The counterfactual condition $\Box_{\gamma'} \varphi$ is true under policies w if φ can be derived from the union of w and γ' . The rest of the connectives are standard. We say that a formula φ is *valid* if $\forall w, w \Vdash \varphi$. We write $\Vdash \varphi$ to mean that φ is valid.

It is important to note that, while Horn clauses may be read as reverse implication, $q :- p$ is not logically equivalent to $p \rightarrow q$. Consider the following example. Let $\gamma = \top$. Then, $\gamma \Vdash p \rightarrow q$ since $p \rightarrow q$ is logically equivalent to $\neg p \vee q$ and $\gamma \not\Vdash p$. But, $\gamma \not\Vdash \Box_p q$ because $\top, p \not\Vdash q$. The counterfactual statement $\Box_p q$ indicates a stronger relation between p and q than the material implication $p \rightarrow q$, so they cannot be used interchangeably.

3.2 Formal Definitions of Probing Attack

Probing attacks can be formally defined based on the semantics shown in Figure 2. A *probe* π is a formula of the form $\Box_{\gamma'} \varphi$, where φ is a \Box -free formula. $\bar{\gamma}'$ is called *the probe credential set* and φ is called *the probe query*. An observation of a probe π under a set of policies w_0 is either π if $w_0 \Vdash \pi$, or $\neg\pi$ otherwise. In our example, $U.\text{canAccessGym}(A)$ is a probe query and a probe credential is $A.\text{registerForGym} :- U.\text{lowGPA}(B)$. A probing attack on w_0 consisting of probes $\{\pi_1, \dots, \pi_n\}$ is the conjunction of the observations of

$\pi_i \in \{\pi_1, \dots, \pi_n\}$ under w_0 . We treat trust management systems as black boxes, so w_0 is not visible to an attacker, only the observations under w_0 . In our example, w_0 is the set of U 's policies, including γ_0 defined above. The probe is the conjunction of φ_1 and φ_2 . If φ is a probing attack on w_0 , then the attacker knows that $w_0 \Vdash \varphi$. If some property φ' holds in all w such that $w \Vdash \varphi$, then the attacker knows for sure that φ' holds under w_0 . We say that φ' is *detectable*.

Definition 1 (Detectability). A formula φ' is detectable in a probing attack φ iff $\forall w$ s.t. $w \Vdash \varphi$, $w \Vdash \varphi'$

Conversely, if there exists some policy w' such that φ holds in w' , but $w' \not\Vdash \varphi'$, then attacker cannot be certain that φ' holds given the results of the probes. Thus, φ' is said to be opaque.

Definition 2 (Opacity). A formula φ' is opaque in a probing attack φ iff it is not detectable in φ , or equivalently, $\exists w_1, w_1 \Vdash \varphi$ and $w_1 \not\Vdash \varphi'$ and $\exists w_2, w_2 \Vdash \varphi$ and $w_2 \not\Vdash \varphi'$

Note that our definition of Opacity differs slightly from that of Becker et al. We made explicit that there exist two policy contexts that produce the same observation for the probe φ , but differ in whether φ' is observable.

Lemma 3 (Probing attacks). A formula φ' is detectable in a probing attack φ iff $\forall w, w \Vdash \varphi \rightarrow \varphi'$

In our example, $\forall w, w \Vdash \varphi_1 \wedge \varphi_2$ implies $w \Vdash U.\text{lowGPA}(B)$. Therefore, $U.\text{lowGPA}(B)$ is detectable in a probing attack consisting of φ_1 and φ_2 (see Section 3.1).

Consider a similar probing attack on a different policy, γ'_0 , also not visible to A, which states that a person can access the gym if they have registered for the gym and paid their fees. The two probes φ_1 and φ'_2 are shown below.

$$\begin{aligned} \gamma'_0 &= U.\text{canAccessGym}(x) :- (x.\text{registerForGym}, x.\text{paidFees}) \\ \varphi_1 &= \neg U.\text{canAccessGym}(A) \\ \varphi'_2 &= \Box_{A.\text{registerForGym} :- U.\text{lowGPA}(B)} \neg U.\text{canAccessGym}(A) \end{aligned}$$

In this example, A sees that they cannot access the gym (φ_1), so they input the credentials: "if B has a low GPA, A registers for the gym" and observes that they still cannot access the gym (φ'_2). From this information, A may want to conclude that B does not have a low GPA ($\neg U.\text{lowGPA}(B)$), but it is easy to construct a w_1 and w_2 to show that $\neg U.\text{lowGPA}(B)$ is opaque given the above two probes. If we let $w_1 = \gamma'_0$, we find that $w_1 \Vdash \varphi_1 \wedge \varphi'_2$ and $w_1 \Vdash \neg U.\text{lowGPA}(B)$. If we let $w_2 = (U.\text{lowGPA}(B), \gamma'_0)$, then, $w_2 \Vdash \varphi_1 \wedge \varphi'_2$, but $w_2 \not\Vdash \neg U.\text{lowGPA}(B)$, therefore $\neg U.\text{lowGPA}(B)$ is opaque in the probing attack $\varphi_1 \wedge \varphi'_2$.

3.3 Proof System

Checking whether a probing attack can detect a secret requires us to check $\forall w, w \Vdash \varphi \rightarrow \varphi'$. This is inconvenient, as we need to examine all possible w . A proof system with derivation rules is helpful in such situations. Becker et al. [5] provides an axiomatization of counterfactual reasoning for Horn clauses (summarized in Figure 3). Here, $\vdash \varphi$ denotes φ is provable by those axioms and proof rules. It has been shown that $\vdash \varphi$ iff $\Vdash \varphi$. Therefore, showing φ' is detectable in probing attack φ is reduced to constructing a proof that $\vdash \varphi \rightarrow \varphi'$.

Next, we explain the axioms and proof rules in more detail. Axioms (C11)-(C13) and Modus Ponens (MP) are from classical propositional logic.

Axioms

$\vdash \varphi \rightarrow \varphi' \rightarrow \varphi$	(C1)
$\vdash (\varphi \rightarrow \varphi' \rightarrow \varphi'') \rightarrow (\varphi \rightarrow \varphi') \rightarrow \varphi \rightarrow \varphi''$	(C12)
$\vdash (\neg\varphi \rightarrow \neg\varphi') \rightarrow \varphi' \rightarrow \varphi$	(C13)
$\vdash \Box_Y (\varphi \rightarrow \varphi') \rightarrow \Box_Y \varphi \rightarrow \Box_Y \varphi'$	(K)
$\vdash \Box_Y \gamma$	(C1)
$\vdash \Box_Y \varphi \rightarrow \gamma \rightarrow \varphi$	(C2)
$\vdash \Box_{(p :- \vec{q})} \varphi \rightarrow (\vec{q} \rightarrow p) \rightarrow \varphi$	(Dlog)
(provided φ is \Box -free)	
$\vdash \Box_Y \neg\varphi \leftrightarrow \neg\Box_Y \varphi$	(Fun)
$\vdash \Box_{Y \wedge Y'} \varphi \leftrightarrow \Box_Y \Box_{Y'} \varphi$	(Perm)

Proof rules

If $\vdash \varphi$ and $\vdash \varphi \rightarrow \varphi'$, then $\vdash \varphi'$	(MP)
If $\vdash \varphi$ then $\vdash \Box_Y \varphi$	(N)
If $\vdash \gamma \rightarrow \gamma'$ and φ is \neg free then $\vdash \Box_{\gamma'} \varphi \rightarrow \Box_Y \varphi$	(Mon)

Figure 3: Axiomatized Proof System

Axiom (K) states that Modus Ponens is preserved under counterfactual conditions. Rule (N) states that if φ is true, then φ is true given additional policies.

Axiom (C1) is a basic requirement for counterfactuals, that if γ is the case, then γ would hold. Axiom (C2) states that counterfactual conditionals are stronger than material implication.

Axiom (Dlog) is the key counterfactual argument for Horn clauses. It captures the essence of Horn clause reasoning. The left hand side of the implication means that “ φ would hold in the policy if the clause $p :- \vec{q}$ were submitted”. The right hand side of the implication, can be expanded as $(\vec{q} \wedge \neg p) \vee \varphi$. This means that the left hand side holds only if: either (1) φ holds in the policy anyway, even without submitting $p :- \vec{q}$, or (2) the clause must be crucial for making φ true. This is only possible if the body of the clause (\vec{q}) are all derivable in the policy, and furthermore that p does not already hold in the policy. Otherwise, the clause could not possibly be crucial.

Axiom (Dlog) only holds for \Box -free φ . Intuitively, if additional clauses can be submitted inside φ , there is no guarantee that the clause on the outermost level is used in the derivation. A counterexample is described in detail in [5]. Briefly, the counterexample leads to $\vdash (p \rightarrow q) \rightarrow \Box_p q$, which is not valid. For $\Box_p q$ to be true in a policy context w , it requires either q to already be true, or both p and q to be true. On the other hand, $p \rightarrow q$ is true if either q is true, or $\neg p$ is true. Note that $(p \rightarrow q)$ and $\Box_p q$ are not logically equivalent to each other. See the end of Section 3.1 for details.

Axiom (Fun) allows negation and the modal operator to commute with each other: if after submitting the policy γ , φ is not derivable, then it is not the case that submitting the policy γ will make φ derivable, and vice versa.

Axiom (Perm) states that submitting policies one by one is equivalent to submitting them all at once.

The proof rule (Mon) states that if φ can be derived given a set of policies γ' , then φ can be derived given a set of logically stronger policies γ . This is only true if φ does not contain negation. Consider an example where $\varphi = \neg p$. Even though $\vdash q \wedge p \rightarrow q$, the formula $\Box_q \neg p \rightarrow \Box_{q,p} \neg p$ is not valid.

The proof system defined in Figure 3 can be used to prove that a probing attack is useful for the attacker. For example, a proof of $\vdash \varphi_1 \wedge \varphi_2 \rightarrow U.\text{lowGPA}(B)$ can be constructed.

The main property of the proof system is that it is sound and complete with regard to semantics. In addition to the derivability semantics, Becker et al. also defined Kripke semantics, as is common for modal logic. The axiomatized proof system is sound and complete with regard to both models, where $\models_{TM} \varphi$ denotes that φ is valid in the Kripke model.

THEOREM 3.1. $\models \varphi \leftrightarrow \models_{TM} \varphi \leftrightarrow \vdash \varphi$.

4 SEQUENT CALCULUS

Axiomatized systems are inconvenient for automated theorem proving and for studying meta-properties of the logic system. For instance, Becker et al. have to prove a series of complex lemmas to manipulate formula contexts. We develop a Gentzen-style sequent calculus for counterfactual reasoning of Horn clauses. A left and a right rule is defined for each connective. The left rule, read from the bottom up, shows how to deconstruct a connective on the left and the right rule, read from the top down, shows how to construct a connective on the right. Derivation rules for each of the logical connectives are defined independently from each other. The sequent calculus rules provide inductive principles to prove meta-properties of the proof system.

Our formalism is inspired by the formalization of classical modal S5 by Murphy et al. [28]. This formalization is a hybrid logic where elements in the Kripke semantics are present in the sequent calculus rules. We write $\varphi @ w$ to denote the counterfactual condition that if w were submitted, then φ would have been true. Here, a world, w , corresponds to a set of policies that are true in that world. The relation between worlds are also made explicit in the proof rules.

4.1 Syntax

Formulas φ and policies γ are the same as in the previous section. We define additional constructs for our sequent calculus below.

<i>Judgment</i>	$J ::= \varphi @ w$
<i>Contexts</i>	$\Delta, \Gamma ::= \cdot \mid \Delta, J$
<i>Positive Forms</i>	$\varphi^+ ::= \gamma^+ \mid \varphi^+ \wedge \varphi'^+ \mid \varphi^+ \vee \varphi'^+$ $\mid \Box_Y \varphi^+ \mid \varphi^- \rightarrow \varphi^+ \mid \neg\varphi^-$
<i>Negative Forms</i>	$\varphi^- ::= \perp \mid \varphi^- \wedge \varphi'^- \mid \varphi^- \vee \varphi'^-$ $\mid \Box_Y \varphi^- \mid \neg\varphi^+ \mid \varphi^+ \rightarrow \varphi^-$

We define positive and negative formulas. Intuitively, positive formulas are the ones that cannot be made false by strengthening the policies and negative formulas are the ones that cannot be made false by weakening the policies. A judgment, denoted J , is a counterfactual condition which states that if w were submitted, then the formula φ is derivable. We write Γ and Δ to denote the contexts that contain a set of counterfactual conditional judgments.

4.2 Derivation Rules

The main judgment of our logic is: $\Gamma \vdash_{\Sigma} \Delta$. This judgement means that if all of the judgments in Γ are true, then one of the judgments in Δ is true. This follows the sequent calculus for classical logic. The context Σ is a finite set of atomic propositions which includes all propositions relevant to the application. For instance, given a trust

$$\begin{array}{c}
\frac{\text{symbols}(\Gamma, p @ w, \Delta) \subseteq \Sigma}{\Gamma, p @ w \vdash_{\Sigma} p @ w, \Delta} \text{INIT} \\
\\
\frac{w \Vdash_d p \quad \Gamma, p @ w \vdash_{\Sigma} \Delta}{\Gamma \vdash_{\Sigma} \Delta} \text{POLCUT} \\
\\
\frac{\Gamma, p @ (w, w', q), q @ w, p @ (w, w') \vdash_{\Sigma} \Delta}{\Gamma, p @ (w, w', q), q @ w \vdash_{\Sigma} \Delta} \text{POLL-W-AT} \\
\\
\frac{\Gamma, p @ (w, w', \Box_{\bar{a}} b), b @ (w, \bar{a}), p @ (w, w') \vdash_{\Sigma} \Delta}{\Gamma, p @ (w, w', \Box_{\bar{a}} b), b @ (w, \bar{a}) \vdash_{\Sigma} \Delta} \text{POLL-W-}\square \\
\\
\frac{\Gamma, s @ (q :- p_1, \dots, p_n, w), s @ w \vdash_{\Sigma} \Delta \quad \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ (q, w), \quad p_1 @ w, \dots, p_n @ w \vdash_{\Sigma} \Delta}{\Gamma, s @ (q :- p_1, \dots, p_n, w) \vdash_{\Sigma} \Delta} \text{DLOG} \\
\\
\frac{\Gamma, p @ w_1, p @ w_2 \vdash_{\Sigma} \Delta \quad \wedge(w_2) @ \top \vdash_{\Sigma} \wedge(w_1) @ \top}{\Gamma, p @ w_1 \vdash_{\Sigma} \Delta} \text{CONV} \\
\\
\frac{\Gamma \vdash_{\Sigma} \varphi @ (w, \bar{y}), \Box_{\gamma} \varphi @ w, \Delta}{\Gamma \vdash_{\Sigma} \Box_{\gamma} \varphi @ w, \Delta} \square R \\
\\
\frac{\Gamma, \Box_{\gamma} \varphi @ w, \varphi @ (w, \bar{y}) \vdash_{\Sigma} \Delta}{\Gamma, \Box_{\gamma} \varphi @ w \vdash_{\Sigma} \Delta} \square L
\end{array}$$

Figure 4: Sequent calculus for policies

management system, Σ would be every atomic proposition in that system, which is finite. We can prove as an invariant of our sequent rules that if $\Gamma \vdash_{\Sigma} \Delta$, then the atomic propositions in context Γ and Δ are a subset of Σ .

We write $\text{symbols}(J)$ to denote the set of atomic propositions in J . We lift it to the logical context and write $\text{symbols}(\Gamma)$ to denote the set of atomic propositions in Γ .

The sequent rules for the connectives from propositional logic are straightforward. They are shown in Figure 5. We omit Σ when it is clear from the context. We explain rules directly related to counterfactual reasoning (Figure 4).

Rule **INIT** states that if we assume that an atomic predicate p can be derived from policies in w , then we can conclude the same. The side condition $\text{symbols}(\Gamma, p @ w, \Delta) \subseteq \Sigma$ ensures that the proof does not include propositions that are not relevant to the system. **POLCUT**, **POLL-W-AT**, **POLL-W- \square** , **DLOG**, and **CONV** resemble left rules, so they are read from the bottom up.

Rule **POLCUT** connects the derivability of a proposition from a set of Horn clauses to counterfactual reasoning. If p is derivable from a policy w , then it must be the case that the counterfactual condition, “if w were submitted, then p is derivable” is true. Therefore, it can be added to the antecedents. While it is not included as a condition here, it will be enforced later that the propositions in p and w are confined to atomic propositions from the set Σ .

$$\begin{array}{c}
\frac{\Gamma, \varphi_1 @ w \vdash_{\Sigma} \varphi_2 @ w, \varphi_1 \rightarrow \varphi_2 @ w, \Delta}{\Gamma \vdash_{\Sigma} \varphi_1 \rightarrow \varphi_2 @ w, \Delta} \rightarrow R \\
\\
\frac{\Gamma, \varphi_1 \rightarrow \varphi_2 @ w \vdash_{\Sigma} \varphi_1 @ w, \Delta \quad \Gamma, \varphi_2 @ w, \varphi_1 \rightarrow \varphi_2 @ w \vdash_{\Sigma} \Delta}{\Gamma, \varphi_1 \rightarrow \varphi_2 @ w \vdash_{\Sigma} \Delta} \rightarrow L \\
\\
\frac{\Gamma \vdash_{\Sigma} \varphi_1 @ w, \varphi_1 \vee \varphi_2 @ w, \Delta}{\Gamma \vdash_{\Sigma} \varphi_1 \vee \varphi_2 @ w, \Delta} \vee R1 \\
\\
\frac{\Gamma \vdash_{\Sigma} \varphi_2 @ w, \varphi_1 \vee \varphi_2 @ w, \Delta}{\Gamma \vdash_{\Sigma} \varphi_1 \vee \varphi_2 @ w, \Delta} \vee R2 \\
\\
\frac{\Gamma, \varphi_1 \vee \varphi_2 @ w, \varphi_1 @ w \vdash_{\Sigma} \Delta \quad \Gamma, \varphi_1 \vee \varphi_2 @ w, \varphi_2 @ w \vdash_{\Sigma} \Delta}{\Gamma, \varphi_1 \vee \varphi_2 @ w \vdash_{\Sigma} \Delta} \vee L \\
\\
\frac{\Gamma \vdash_{\Sigma} \varphi_1 @ w, \varphi_1 \wedge \varphi_2 @ w, \Delta \quad \Gamma \vdash_{\Sigma} \varphi_2 @ w, \varphi_1 \wedge \varphi_2 @ w, \Delta}{\Gamma \vdash_{\Sigma} \varphi_1 \wedge \varphi_2 @ w, \Delta} \wedge R \\
\\
\frac{\Gamma, \varphi_1 \wedge \varphi_2 @ w, \varphi_1 @ w \vdash_{\Sigma} \Delta}{\Gamma, \varphi_1 \wedge \varphi_2 @ w \vdash_{\Sigma} \Delta} \wedge L1 \\
\\
\frac{\Gamma, \varphi_1 \wedge \varphi_2 @ w, \varphi_2 @ w \vdash_{\Sigma} \Delta}{\Gamma, \varphi_1 \wedge \varphi_2 @ w \vdash_{\Sigma} \Delta} \wedge L2
\end{array}$$

$$\frac{\Gamma, \neg \varphi @ w \vdash_{\Sigma} \varphi @ w, \Delta}{\Gamma, \neg \varphi @ w \vdash_{\Sigma} \Delta} \neg L \quad \frac{\Gamma, \varphi @ w \vdash_{\Sigma} \neg \varphi @ w, \Delta}{\Gamma \vdash_{\Sigma} \neg \varphi @ w, \Delta} \neg R$$

$$\frac{\text{symbols}(\Gamma, \top @ w, \Delta) \subseteq \Sigma}{\Gamma \vdash_{\Sigma} \top @ w, \Delta} \top R$$

$$\frac{\text{symbols}(\Gamma, \perp @ w, \Delta) \subseteq \Sigma}{\Gamma, \perp @ w \vdash_{\Sigma} \Delta} \perp L$$

Figure 5: Sequent calculus for classical propositional logic

POLL-W-AT and **POLL-W- \square** are rules for simplifying policies based on the derivability of Horn clauses. **POLL-W-AT** says that if we assume p is true given w, w' , and q , and we assume that q is true given w , then it is safe to assume that p is true given only w and w' ; q is redundant. **POLL-W- \square** is similar except that the redundant policy is of form $\Box_{\bar{a}} b$.

The **DLOG** rule introduces counterfactual reasoning of Horn clauses into the sequent rules. It says that if we know that an atomic proposition s is true given the policies in w and $q :- p_1, \dots, p_n$, then it must be the case that either s is already true given w ; or if rule $q :- p_1, \dots, p_n$ were not present, then s could be derived. The latter part is reflected in the last premise, where the body propositions of the rule are derivable in w , and s is derivable under w and the head of the rule (q).

$$\begin{array}{c}
\frac{\dots, c @ (w, a, b), a @ w, c @ (w, b) \vdash c @ (w, b) \dots}{\dots, c @ (w, a, b), a @ w \vdash c @ (w, b), \Box_b c @ w, \dots} \text{INIT} \\
\frac{\dots, c @ (w, a, b), a @ w \vdash c @ (w, b), \Box_b c @ w, \dots}{\dots, \Box_{a \wedge b} c @ w, c @ (w, a, b), a @ w \vdash \Box_b c @ w, \dots} \text{PoL-W-AT} \\
\frac{\dots, \Box_{a \wedge b} c @ w, c @ (w, a, b), a @ w \vdash \Box_b c @ w, \dots}{\dots, \Box_{a \wedge b} c @ w, a @ w \vdash \Box_b c @ w, \neg a @ w} \text{\(\Box\)R} \\
\frac{\dots, \Box_{a \wedge b} c @ w, a @ w \vdash \Box_b c @ w, \neg a @ w}{\neg \Box_b c @ w, \Box_{a \wedge b} c @ w \vdash \Box_b c @ w, \neg a @ w} \text{\(\Box\)L} \\
\frac{\neg \Box_b c @ w, \Box_{a \wedge b} c @ w \vdash \Box_b c @ w, \neg a @ w}{\neg \Box_b c @ w, \Box_{a \wedge b} c @ w \vdash \neg a @ w} \text{\(\neg\)R} \\
\frac{\neg \Box_b c @ w, \Box_{a \wedge b} c @ w \vdash \neg a @ w}{\neg \Box_b c @ w, \Box_{a \wedge b} c @ w \vdash \neg a @ w} \text{\(\neg\)L}
\end{array}$$

Figure 6: Simple probing attack proof

The next rule allows us to change the world (policies) under which a formula is derived. CONV makes explicit the relation between worlds. If we assume that p is derivable after submitting policies w_1 ($p @ w_1$), then it is safe to assume $p @ w_2$ given that w_1 is weaker than w_2 . More simply, this rule states that if p is derivable given w_1 and w_1 is derivable from w_2 , then p should also be derivable from w_2 . Similar to the POLCUT rule, the newly introduced world w_2 can only be composed of symbols from Σ .

The left and right rules for the \Box connective simply add the policy γ from the index of the modal operator \Box to the world under which the formula is queried. The right rule states that if φ is true given w and γ , then $\Box_\gamma \varphi$ is true given only w . The left rule says that to use $\Box_\gamma \varphi @ w$ in the context, decompose it and use the fact that φ is true under w and $\bar{\gamma}$. This rule is closely related to the Kripke semantics of the \Box modality. The Kripke semantics require φ to be true in all worlds that are reachable from w with distance γ . In this rule, we only select one such world, which is $(w, \bar{\gamma})$. As a result, we compensate with our conversion rule (see Section 6).

4.3 Example Derivations

We show in Figure 6 the proof for the counterfactual reasoning: $\vdash \neg \Box_b c \wedge \Box_{a \wedge b} c \rightarrow \neg a$. The probes tell us the following. If we submit b , c cannot be derived. However, if we submit both a and b , c can be derived. From the probes, we can conclude that in the current system, a cannot be true; otherwise, submitting b would have allowed the derivation of c . Reading from the bottom up, first $\neg L$, $\neg R$, $\Box L$, and $\Box R$ rules are used to reach the subgoal: if c is derivable given w , a , and b , and a is derivable from w , then c is derivable from w and b . Intuitively, this is correct because w subsumes a . This subgoal can be discharged by applying the POL-W-AT rule. We complete the proof by applying the INIT rule.

Next, we show the proof of the probing attack discussed in Section 1, formalized in Section 3.1. We use the following abbreviation: $A_a = U.\text{canAccessGym}(A)$, $G_b = U.\text{lowGPA}(B)$, and $R_a = A.\text{registerForGym}$. We write $\mathcal{E} ::$ to label derivations so they can be referred to later.

$$\frac{\mathcal{E}_1 :: \neg A_a @ \top, \dots, A_a @ \top \vdash G_b @ \top \quad \mathcal{E}_2 :: \dots, G_b @ \top \vdash G_b @ \top}{\neg A_a @ \top, \dots, A_a @ (\top, R_a :- G_b) \vdash G_b @ \top} \text{DLog} \\
\frac{\neg A_a @ \top, \dots, A_a @ (\top, R_a :- G_b) \vdash G_b @ \top}{\neg A_a @ \top, \Box_{R_a :- G_b} A_a @ \top \vdash G_b @ \top} \text{\(\Box\)L}$$

Once we apply the DLOG rule, we only need to prove that if A currently has access to the gym then B has a low GPA (\mathcal{E}_1) and that

if B has a low GPA, then B has a low GPA (\mathcal{E}_2). \mathcal{E}_1 contains a contradiction on the left, because our first probe says that currently A does not have access to the gym. We can apply $\neg L$ rule to discharge it. \mathcal{E}_2 can be proven by applying the INIT rule.

The last example proves a more complex probing attack from [5]. Consider the following formula:

$$s @ a, \neg s @ (a :- b), s @ (a :- b, b :- x) \vdash x @ \top$$

x is a sensitive credential that should not be leaked to the attacker. On the left of the turnstile are three probes by the attacker. The proof shows that x is detectable under these probes. Informally, from the first probe, the attacker knows that if a is submitted, then s is true. The second probe shows that if $(a :- b)$ is submitted, then s cannot be derived. From these two probes, the attacker knows that b cannot be true. Otherwise, the second probe would have shown that s is derivable after $(a :- b)$ is submitted. The third probe shows that s is derivable if two policies: $(a :- b)$ and $(b :- x)$ are submitted. We already know from the second probe that $(a :- b)$ alone cannot make s derivable, so $(b :- x)$ is crucial in deriving s . Since we already know that b is not true in the current context, it must be the case that the body of $(b :- x)$ is true, which is the secret that the attacker tries to guess. We explain the proofs using our sequent calculus rules, reading from the bottom up (Figure 7). We mark formulas used on the rules in red.

To avoid using Cut, our proof analyzes the last probe first. We apply DLOG rule on the last probe and deconstruct the rule $(b :- x)$. The first subgoal \mathcal{D} contains a contradiction. This basically states that there is no way s can be derived only from $(a :- b)$. The second subgoal \mathcal{E} analyzes the situation where x is true under the policy $(a :- b)$ and s is derivable assuming b and $(a :- b)$ are derivable. To prove \mathcal{E} , we apply DLOG rule and analyze the assumption: x is true under the policy $(a :- b)$. There are two cases: either x is true without using $a :- b$, in which case we succeed in proving that $x @ \top$; or x is true given a , and b is true. The second case is denoted \mathcal{E}' . The proof of \mathcal{E}' essentially derives a contradiction: s can be derived from $a :- b$. This is by using the $\neg L$ first, then applying the POL-W-AT rule which shows that s must be derivable from $a :- b$, because b is true and s can be derived from b and $a :- b$.

In the end, we cover all the cases, and show that x is detectable under the three probes.

5 METATHEORY

We prove that our sequent calculus has cut elimination, is consistent, and is sound and complete with regard to the axiomatized

$$\begin{array}{l}
\mathcal{D} :: s @ a, \neg s @ (a :- b), s @ (a :- b, b :- x), s @ (a :- b) \vdash x @ \top \\
\mathcal{E} :: s @ a, \neg s @ (a :- b), s @ (a :- b, b :- x), s @ (b, a :- b), x @ (a :- b) \vdash x @ \top \\
\hline
s @ a, \neg s @ a :- b, s @ (a :- b, b :- x) \vdash x @ \top \quad \text{DLog} \\
\\
\mathcal{D} \quad \text{can be proving using } \neg\text{-L and INIT} \\
\\
\mathcal{E} = \frac{\cdots, x @ \top \vdash x @ \top \quad \mathcal{E}' :: \cdots, \neg s @ (a :- b), s @ (b, a :- b), x @ (a, \top), b @ \top \vdash x @ \top}{s @ a, \neg s @ (a :- b), s @ (b, a :- b), x @ a :- b \vdash x @ \top} \text{DLog} \\
\\
\mathcal{E}' = \frac{\frac{\cdots, s @ (b, a :- b), s @ (a :- b), b @ \top \vdash s @ (a :- b), x @ \top}{\cdots, s @ (b, a :- b), b @ \top \vdash s @ (a :- b), x @ \top} \text{PoLL-W-At} \quad \text{INIT}}{\cdots, \neg s @ (a :- b), s @ (b, a :- b), b @ \top \vdash x @ \top} \neg\text{-L}
\end{array}$$

Figure 7: Complex probing attack proof

formulation in Section 3. At the end of this section, we prove correct a defense against probing attacks. The proof inducts over the structure of the derivation, which is a huge benefit of sequent calculus compared to axiomatized systems.

5.1 Admissibility of Cut (Cut Elimination)

One of the most important properties of a sequent calculus is the admissibility of cut (Theorem 4). We write $\mathcal{D} ::$, $\mathcal{E} ::$, or $\mathcal{F} ::$ to label derivations so they can be referred to later. We can prove the admissibility of cut using similar techniques outlined by Pfenning [33].

Theorem 4 (Admissibility of Cut).

If $\mathcal{D} :: \Gamma \vdash \varphi @ w, \Delta$ and $\mathcal{E} :: \Gamma, \varphi @ w \vdash \Delta$, then $\exists \mathcal{F} :: \Gamma \vdash \Delta$

Proof (sketch): Following Pfenning [33], we refer to a formula that is introduced on the left or right as a principal formula. For the INIT rule, the principal formula is the formula appearing on both sides. For the \top R and \perp L rules, the principal formula is the $\top @ w$ formula appearing on the right and the $\perp @ w$ formula appearing on the left, respectively. We refer to all other formulas as side formulas.

The proof is by nested induction, first over the structure of the cut formula, φ , and then over the structures of \mathcal{D} and \mathcal{E} . All proof cases fall into 4 categories (many of these cases overlap), more details may be found in Appendix C.

Base cases: Either \mathcal{D} or \mathcal{E} ends in INIT, \top R, or \perp L

- Either \mathcal{D} or \mathcal{E} ends in INIT and the other derivation is arbitrary. The cut formula is $p @ w$. Applying the contraction lemma on the arbitrary derivation will result in \mathcal{F} .
- \mathcal{D} is an arbitrary derivation and \mathcal{E} ends in \top R. Then, the cut formula is an arbitrary $\varphi @ w$. This case is covered by side cases on \mathcal{E} .
- \mathcal{D} ends in \top R and \mathcal{E} is an arbitrary derivation. Then, the cut formula is $\top @ w$. This case is covered by side cases on \mathcal{E} .
- \mathcal{D} ends in \perp L and \mathcal{E} is an arbitrary derivation. Then, the cut formula is an arbitrary $\varphi @ w$. This case is covered by side cases on \mathcal{D} .
- \mathcal{D} is an arbitrary derivation and \mathcal{E} ends in \perp L. Then, the cut formula is $\perp @ w$. This case is covered by side cases on \mathcal{D} .

Principal cases: The cut formula is the principal formula in both \mathcal{D} and \mathcal{E} . \mathcal{D} and \mathcal{E} end in matching left and right rules of the same connective. (1) Apply the induction hypothesis on the cut formula, \mathcal{D} , and each of the subderivations of \mathcal{E} . (2) Apply the induction hypothesis on the cut formula, each of the subderivations of \mathcal{D} , and \mathcal{E} . Applying the induction hypothesis on a subformula of the cut formula and the results of (1) and (2) gives \mathcal{F} .

Special cases: \mathcal{E} ends in PoLCut, PoLL-W-At, PoLL-W- \square , DLog, or Conv and the cut formula is one of the formulas used by the rule. These cuts may be of interest since they involve our new rules, but they do not actually require separate proof cases. All of these types of cuts are covered by other cases, outlined below.

- \mathcal{D} ends in INIT. This case is covered by the base cases where \mathcal{D} is INIT and \mathcal{E} is an arbitrary derivation.
- The cut formula is a side formula in \mathcal{D} . This case is covered by the side cases on \mathcal{D} .

Side cases on \mathcal{D} : The cut formula is a side formula in \mathcal{D} . There is a case for \mathcal{D} ending in every rule. \mathcal{E} is always an arbitrary derivation.

- \mathcal{D} ends in INIT, \top R, or \perp L: We can directly apply the rule which \mathcal{D} ends in to obtain \mathcal{F} .
- Otherwise: We can apply the rule which \mathcal{D} ends in on the result of applying the induction hypothesis to the cut formula, the subderivations of \mathcal{D} , and \mathcal{E} , to obtain \mathcal{F} .

Side cases on \mathcal{E} : The cut formula is a side formula in \mathcal{E} . There is a case for \mathcal{E} ending in every rule. \mathcal{D} is always an arbitrary derivation.

- \mathcal{E} ends in INIT, \top R, or \perp L: We can directly apply the rule which \mathcal{E} ends in to obtain \mathcal{F} .
- Otherwise: We can apply the rule which \mathcal{E} ends in on the result of applying the induction hypothesis to the cut formula, \mathcal{D} , and the subderivations of \mathcal{E} , to obtain \mathcal{F} .

□

5.2 Consistency

Normally, with the admissibility of cut, the consistency of the logic becomes a trivial proof by showing that there is no rule to derive

$\cdot \vdash \perp @ w$. However, for us, the **POLCUT** rule could be used. To prove consistency, we need to generalize the statement and induct over the size of the derivation. Theorem 5 allows any counterfactual condition of form $p @ w$ in Γ . The ordinary consistency property follows directly from it.

Theorem 5 (Consistency).

$\forall n, \#E, \#w, \#\Gamma$ such that $E :: \Gamma \vdash \perp @ w$ and $|\mathcal{E}| \leq n$ where $\forall J \in \Gamma, J = p @ w$.

5.3 Soundness

For soundness, we need to show that if φ is derivable using our sequent rules, it is also derivable using the proof system in Figure 3. Instead of proving this directly, we prove the soundness of our sequent calculus with regard to the formula semantics (Theorem 6). We write $\bigwedge \Gamma$ to denote the conjunction of the counterfactuals that correspond to judgments in Γ . That is, the counterfactual that corresponds to the judgment $\varphi @ w$ is $\square_w \varphi$. Similarly, $\bigvee \Delta$ is the disjunction of the counterfactuals that corresponds to judgments in Δ . In combination with Theorem 3.1, Theorem 6 allows us to connect formulas derivable in our sequent rules to those derivable using the axioms.

Theorem 6 (Soundness).

If $E :: \Gamma \vdash \Delta$, then $\forall w, w \Vdash \neg(\bigwedge \Gamma) \vee \bigvee \Delta$.

The proof is by induction on the structure of \mathcal{E} . Most cases are straightforward. The most interesting case is where \mathcal{E} ends in **CONV**. The I.H. on the first premise gives us

$$\mathcal{E}'_1 :: w \Vdash \neg \bigvee \Gamma \vee \square_{w_1} \neg p \vee \square_{w_2} \neg p \vee \bigvee \Delta$$

but we want to show

$$w \Vdash \neg \bigvee \Gamma \vee \square_{w_1} \neg p \vee \bigvee \Delta$$

There are several cases to consider here; namely, when $w \Vdash \neg \bigvee \Gamma$, when $w \Vdash \square_{w_1} \neg p$, and so on. The important case here is when it's not immediately obvious if we can remove $\square_{w_2} \neg p$ from \mathcal{E}'_1 , the case where

$$w \Vdash \square_{w_2} \neg p$$

We need to conclude that $w, w_2 \Vdash \neg p$ implies $w, w_1 \Vdash \neg p$. This follows from our intuition about **CONV** because the last premise of **CONV**, $\bigwedge(w_2) @ \top \vdash \bigwedge(w_1) @ \top$, suggests that w_2 contains stronger policies than w_1 . If p is not derivable from w_2 , then it should not be derivable from w_1 . However, it turns out that the I.H. on the last premise only gives us $\forall w', w' \Vdash \neg(\bigwedge w_2) \vee \bigwedge w_1$, which is not strong enough. Instead, we need $w_2 \Vdash \bigwedge(w_1)$.

In fact, we can prove the following lemma, which states that if a policy γ_2 can be derived from another policy γ_1 using our sequent rules, then γ_2 is stronger than γ_1 under datalog semantics.

Lemma 7. $\forall \gamma_1, \gamma_2$, If $E :: \gamma_1 @ \top \vdash \gamma_2 @ \top$ then $\overline{\gamma_1} \Vdash \gamma_2$.

This lemma confirms that our sequent rules \vdash reason about counterfactual conditions. If, by using our main judgment, we can prove that γ_2 can be derived with only the additional assumption that γ_1 can be derived with no additional policy assumptions, then γ_2 must be derivable from γ_1 .

Proving Lemma 7 is non-trivial. Inducting over the structure of $E :: \gamma_1 @ \top \vdash \gamma_2 @ \top$ proves insufficient because **POLCUT** and

CONV can introduce judgements that our induction hypothesis does not account for. We want to show that we can safely disregard these judgements because they are derivable facts and are made redundant by what we already know. Intuitively, given a sub-derivation of $\gamma_1 @ \top \vdash \gamma_2 @ \top$ of the form $\Gamma, \Gamma_A \vdash \Delta$, where Γ_A contains all the new judgments introduced via **POLCUT** and **CONV**, it is the case that $\bigwedge \Gamma \Vdash \bigwedge \Gamma_A$. If $p @ w$ is introduced by **POLCUT**, we already know $w \Vdash_d p$. If $p @ w_2$ is introduced by **CONV**, we already have $p @ w_1$ in Γ , and w_2 is stronger than w_1 (by I.H.), so $p @ w_2$ can be derived from Γ .

To formally prove this statement, we first define a new calculus containing labels. The labels allow us to track which judgements were introduced by **POLCUT** and **CONV**, as well as track which judgements were derived from those judgements. This will allow us to distinguish judgements of the form $\gamma @ \top$ from those of the form $p @ w$. We assign the label 'A' to judgements that have been added by, or are derived from judgements which have been added by, **POLCUT**⁺ or **CONV**⁺. All other judgements are labeled with 'O'.

For example, consider the following rule from our labeled sequent calculus:

$$\frac{\text{symbols}(w_2) \subseteq \Sigma \quad \Gamma^+, (p @ w_1)_x, (p @ w_2)_A \vdash \Delta}{(\bigwedge(w_2) @ \top)_o \vdash \bigwedge(w_1) @ \top} \text{CONV}^+$$

$p @ w_2$ is a judgement that is being added by **CONV**, so it is labeled with an 'A' to indicate this. We know that $\bigwedge(w_2)$ is derivable after submitting \top , so it can safely be labeled with an 'O'. The remaining rules, as well as details about the syntax and notation, may be found in Appendix A.

Next, we prove that derivations using the original sequent rules can be annotated, resulting in derivations in the labeled sequent rules. Finally, we prove a more general lemma, which formalizes our intuition of discarding judgments introduced via **POLCUT** and **CONV**, discussed earlier. We show a less precise and simplified version of the lemma below to illustrate how the lemma relates derivations in the marked up sequent rules to formula semantics.

Lemma 8. If $\Gamma^+, \Gamma_A^+ \vdash \Delta$ and $\bigwedge \Gamma^+ \Vdash \bigwedge \Gamma_A^+$ and judgments in the contexts are can only be form $\gamma @ \top, q @ w$ then $\exists \varphi @ w \in \Delta$ s.t. $\bigwedge \Gamma^+ \Vdash \square_w \varphi$.

In the above lemma Γ_A^+ contains all the judgments that are labeled 'A', which means that it is derived from judgments injected by **POLCUT** and **CONV** rules. The lemma states that if judgments marked by the label A are derivable by counterfactuals in the rest of the context, then one of the counterfactuals in Δ can be derived from the conjunction of counterfactuals in Γ^+ , which contain only judgments decomposed from the judgments in the final conclusion. Lemma 7 is a corollary. The precise lemmas are defined and proven in Appendix A.

5.4 Completeness

For completeness, we prove that for every provable formula φ in the axiomatized system (Figure 3), we can derive $\varphi @ w$ for any w (Theorem 9).

Theorem 9 (Completeness). If $\vdash \varphi$, then $\forall w, \cdot \vdash \varphi @ w$.

Axiom C11

$$\frac{\frac{\frac{}{\varphi_1 @ w, \varphi_2 @ w \vdash_{\Sigma} \varphi_1 @ w, \dots} \text{INIT}}{\varphi_1 @ w \vdash_{\Sigma} \varphi_2 \rightarrow \varphi_1 @ w, \dots} \rightarrow R}{\vdash_{\Sigma} \varphi_1 \rightarrow (\varphi_2 \rightarrow \varphi_1) @ w} \rightarrow R$$

Axiom C12

$$\frac{\frac{\frac{\frac{}{\mathcal{E} :: \varphi_1 \rightarrow (\varphi_2 \rightarrow \varphi_3) @ w, \varphi_1 \rightarrow \varphi_2 @ w, \varphi_1 @ w \vdash_{\Sigma} \varphi_3 @ w, \dots} \rightarrow R}{\varphi_1 \rightarrow (\varphi_2 \rightarrow \varphi_3) @ w, \varphi_1 \rightarrow \varphi_2 @ w \vdash_{\Sigma} \varphi_1 \rightarrow \varphi_3 @ w, \dots} \rightarrow R}{\varphi_1 \rightarrow (\varphi_2 \rightarrow \varphi_3) @ w \vdash_{\Sigma} ((\varphi_1 \rightarrow \varphi_2) \rightarrow (\varphi_1 \rightarrow \varphi_3)) @ w, \dots} \rightarrow R}{\vdash_{\Sigma} (\varphi_1 \rightarrow \varphi_2 \rightarrow \varphi_3) \rightarrow ((\varphi_1 \rightarrow \varphi_2) \rightarrow (\varphi_1 \rightarrow \varphi_3)) @ w} \rightarrow R$$

We complete this derivation by repeatedly applying $\rightarrow L$ and INIT to \mathcal{E} .

Axiom C13

$$\frac{\frac{\frac{}{\mathcal{E} :: \neg\varphi_1 \rightarrow \neg\varphi_2 @ w \vdash_{\Sigma} \neg\varphi_1 @ w, \varphi_2 \rightarrow \varphi_1 @ w, \dots} \rightarrow L}{\mathcal{D} :: \dots, \neg\varphi_2 @ w \vdash_{\Sigma} \varphi_2 \rightarrow \varphi_1 @ w, \dots} \rightarrow L}{\neg\varphi_1 \rightarrow \neg\varphi_2 @ w \vdash_{\Sigma} \varphi_2 \rightarrow \varphi_1 @ w, \dots} \rightarrow R}{\vdash_{\Sigma} (\neg\varphi_1 \rightarrow \neg\varphi_2) \rightarrow (\varphi_2 \rightarrow \varphi_1) @ w} \rightarrow R$$

$$\mathcal{E} = \frac{\frac{\frac{}{\dots, \varphi_2 @ w, \varphi_1 @ w \vdash_{\Sigma} \varphi_1 @ w, \dots} \text{INIT}}{\dots, \varphi_2 @ w \vdash_{\Sigma} \varphi_1 @ w, \neg\varphi_1 @ w, \dots} \rightarrow R}{\neg\varphi_1 \rightarrow \neg\varphi_2 @ w \vdash_{\Sigma} \neg\varphi_1 @ w, \varphi_2 \rightarrow \varphi_1 @ w, \dots} \rightarrow R$$

$$\mathcal{D} = \frac{\frac{\frac{}{\dots, \varphi_2 @ w \vdash_{\Sigma} \varphi_2 @ w, \varphi_1 @ w, \dots} \text{INIT}}{\dots, \neg\varphi_2 @ w, \varphi_2 @ w \vdash_{\Sigma} \varphi_1 @ w, \dots} \rightarrow L}{\dots, \neg\varphi_2 @ w \vdash_{\Sigma} \varphi_2 \rightarrow \varphi_1 @ w, \dots} \rightarrow R$$

Axiom Fun

$$\frac{\frac{\frac{\frac{}{\dots, \varphi @ (w, \bar{\gamma}) \vdash_{\Sigma} \varphi @ (w, \bar{\gamma}), \dots} \text{INIT}}{\square_{\gamma} \varphi @ w \vdash_{\Sigma} \varphi @ (w, \bar{\gamma}), \dots} \rightarrow R}{\square_{\gamma} \varphi @ w \vdash_{\Sigma} \varphi @ (w, \bar{\gamma}), \dots} \rightarrow R}{\dots \vdash_{\Sigma} \varphi @ (w, \bar{\gamma}), \neg \square_{\gamma} \varphi @ w} \rightarrow L}{\dots, \neg \varphi @ (w, \bar{\gamma}) \vdash_{\Sigma} \neg \square_{\gamma} \varphi @ w} \rightarrow L}{\square_{\gamma} \neg \varphi @ w \vdash_{\Sigma} \neg \square_{\gamma} \varphi @ w} \rightarrow L$$

$$\frac{\frac{\frac{\frac{}{\dots, \varphi @ (w, \bar{\gamma}) \vdash_{\Sigma} \varphi @ (w, \bar{\gamma}), \dots} \text{INIT}}{\dots, \varphi @ (w, \bar{\gamma}) \vdash_{\Sigma} \square_{\gamma} \varphi @ w, \dots} \rightarrow L}{\dots, \varphi @ (w, \bar{\gamma}) \vdash_{\Sigma} \square_{\gamma} \varphi @ w, \dots} \rightarrow L}{\neg \square_{\gamma} \varphi @ w, \varphi @ (w, \bar{\gamma}) \vdash_{\Sigma} \dots} \rightarrow R}{\neg \square_{\gamma} \varphi @ w \vdash_{\Sigma} \neg \varphi @ (w, \bar{\gamma}), \dots} \rightarrow R}{\neg \square_{\gamma} \varphi @ w \vdash_{\Sigma} \square_{\gamma} \neg \varphi @ w} \rightarrow R$$

Axiom Perm

$$\frac{\frac{\frac{\frac{}{\dots, \varphi @ (w, \bar{\gamma} \wedge \bar{\gamma}') \vdash_{\Sigma} \varphi @ (w, \bar{\gamma}, \bar{\gamma}'), \dots} \text{INIT}}{\dots, \varphi @ (w, \bar{\gamma} \wedge \bar{\gamma}') \vdash_{\Sigma} \square_{\gamma'} \varphi @ (w, \bar{\gamma}), \dots} \rightarrow R}{\dots, \varphi @ (w, \bar{\gamma} \wedge \bar{\gamma}') \vdash_{\Sigma} \square_{\gamma} \square_{\gamma'} \varphi @ w} \rightarrow R}{\square_{\gamma \wedge \gamma'} \varphi @ w \vdash_{\Sigma} \square_{\gamma} \square_{\gamma'} \varphi @ w} \rightarrow L$$

$$\frac{\frac{\frac{\frac{}{\dots, \varphi @ (w, \bar{\gamma}, \bar{\gamma}') \vdash_{\Sigma} \varphi @ (w, \bar{\gamma} \wedge \bar{\gamma}'), \dots} \text{INIT}}{\dots, \square_{\gamma'} \varphi @ (w, \bar{\gamma}) \vdash_{\Sigma} \varphi @ (w, \bar{\gamma} \wedge \bar{\gamma}'), \dots} \rightarrow L}{\square_{\gamma} \square_{\gamma'} \varphi @ w \vdash_{\Sigma} \varphi @ (w, \bar{\gamma} \wedge \bar{\gamma}'), \dots} \rightarrow L}{\square_{\gamma} \square_{\gamma'} \varphi @ w \vdash_{\Sigma} \square_{\gamma \wedge \gamma'} \varphi @ w} \rightarrow R$$

Axiom K

$$\frac{\frac{\frac{\frac{}{\mathcal{E} :: \dots, \varphi_1 \rightarrow \varphi_2 @ (w, \bar{\gamma}) \vdash_{\Sigma} \varphi_1 @ (w, \bar{\gamma}), \square_{\gamma} \varphi_1 \rightarrow \square_{\gamma} \varphi_2 @ w, \dots} \rightarrow L}{\dots, \varphi_1 \rightarrow \varphi_2 @ (w, \bar{\gamma}) \vdash_{\Sigma} \square_{\gamma} \varphi_1 \rightarrow \square_{\gamma} \varphi_2 @ w, \dots} \rightarrow L}{\square_{\gamma} (\varphi_1 \rightarrow \varphi_2) @ w \vdash_{\Sigma} \square_{\gamma} \varphi_1 \rightarrow \square_{\gamma} \varphi_2 @ w, \dots} \rightarrow L}{\vdash_{\Sigma} \square_{\gamma} (\varphi_1 \rightarrow \varphi_2) \rightarrow \square_{\gamma} \varphi_1 \rightarrow \square_{\gamma} \varphi_2 @ w} \rightarrow R$$

$$\mathcal{E} = \frac{\frac{\frac{\frac{}{\dots, \varphi_1 @ (w, \bar{\gamma}) \vdash_{\Sigma} \varphi_1 @ (w, \bar{\gamma}), \square_{\gamma} \varphi_2 @ w, \dots} \text{INIT}}{\dots, \square_{\gamma} \varphi_1 @ w \vdash_{\Sigma} \varphi_1 @ (w, \bar{\gamma}), \square_{\gamma} \varphi_2 @ w, \dots} \rightarrow R}{\dots, \varphi_1 \rightarrow \varphi_2 @ (w, \bar{\gamma}) \vdash_{\Sigma} \varphi_1 @ (w, \bar{\gamma}), \square_{\gamma} \varphi_1 \rightarrow \square_{\gamma} \varphi_2 @ w, \dots} \rightarrow R$$

$$\mathcal{D} = \frac{\frac{\frac{\frac{}{\dots, \square_{\gamma} \varphi_1 @ w, \varphi_2 @ (w, \bar{\gamma}) \vdash_{\Sigma} \varphi_2 @ (w, \bar{\gamma}), \dots} \text{INIT}}{\dots, \square_{\gamma} \varphi_1 @ w, \varphi_2 @ (w, \bar{\gamma}) \vdash_{\Sigma} \square_{\gamma} \varphi_2 @ w, \dots} \rightarrow R}{\dots, \varphi_2 @ (w, \bar{\gamma}) \vdash_{\Sigma} \square_{\gamma} \varphi_1 \rightarrow \square_{\gamma} \varphi_2 @ w, \dots} \rightarrow R$$

Figure 8: Derivations for the proof of Axioms C11, C12, C13, Fun, Perm, and K

The proof of the completeness theorem shows that all of the axioms can be derived using the sequent calculus rules and that the proof rules are admissible in our sequent calculus. Axioms C11, C12, C13, Fun, Perm, and K can be directly constructed using our sequent

rules. The derivations may be found in Figure 8. The remaining axioms can be shown by induction over the structure of φ and we list the lemmas for these axioms below.

Lemma 10 (Axiom C1). $\forall \gamma, w, \vdash_{\Sigma} \gamma @ (\bar{\gamma}, w)$.

Lemma 11 (Axiom C2). $\forall \varphi, \gamma, w, \vdash_{\Sigma} (\Box_{\gamma} \varphi \rightarrow \gamma \rightarrow \varphi) @ w$.

Lemma 12 (Axiom Dlog). $\forall w, \varphi, p, \vec{q}, \varphi$, if φ is \Box free then $\vdash_{\Sigma} \Box_{(q \rightarrow \vec{p})} \varphi \rightarrow (\vec{p} \rightarrow q) \rightarrow \varphi @ w$

Lemma 13 (Axiom MP). If $\vdash \varphi_1 @ w$ and $\vdash \varphi_1 \rightarrow \varphi_2 @ w$ then $\vdash \varphi_2 @ w$.

Lemma 14 (Axiom N).

$\forall \varphi, w, \gamma$, if $\vdash_{\Sigma} \varphi @ w$, then $\vdash_{\Sigma} \Box_{\gamma} \varphi @ w$

The proof of Axiom Mon follows from a more general lemma for positive and negative formulas and requires additional lemmas to remove unnecessary judgements from the context. The proofs of these lemmas and Axiom Mon can be found in Appendix E.

Lemma 15 (Axiom Mon). If $\vdash_{\Sigma} \gamma_1 \rightarrow \gamma_2 @ \top$ then $\forall w \vdash_{\Sigma} \Box_{\gamma_2} \varphi^+ \rightarrow \Box_{\gamma_1} \varphi^+ @ w$.

Axiom Mon requires \neg -free formulas. Our Lemma 15 is a more general form that allows not only \neg -free formulas, but also other positive formulas.

6 DISCUSSION

Alternative Formalizations In our initial attempt, we tried to use only one policy context and allow Γ and Δ to only contain formulas. The judgment is of the form: $w; \Gamma \vdash \Delta$. The reading of this judgment is that if policies in w are submitted, and all of the formulas in Γ are derivable, then one of the formulas in Δ is derivable. The precise meaning of the judgment with regard to the model is the following: If $w; \Gamma \vdash \Delta$ then $\forall w', w', w \Vdash \bigwedge \Gamma \rightarrow \bigvee \Delta$.

The $\Box R$ and $\Box L$ rules are as follows.

$$\frac{w, \gamma; \Gamma_a, \Gamma \vdash \Delta_b, \Delta \quad \begin{array}{l} \Gamma_a \text{ contains only positive formulas} \\ \Delta_b \text{ contains only negative formulas} \end{array}}{w; \Gamma_a, \Box_{\gamma} \Gamma \vdash \Delta_b, \Box_{\gamma} \Delta} \Box R'$$

$$\frac{w; \Box_{\gamma} \varphi, \Gamma \vdash \gamma, \Delta \quad w; \Box_{\gamma} \varphi, \varphi, \Gamma \vdash \Delta}{w; \Box_{\gamma} \varphi, \Gamma \vdash \Delta} \Box L'$$

We write $\Box_{\gamma} \Gamma$ to denote the formula context resulting from wrapping every formula in Γ with \Box_{γ} . The $\Box R'$ introduces counterfactuals on both sides of the turnstile, except for formulas in contexts Γ_a and Δ_b . The first premise of this rule states that if policies w and γ were submitted, and all the formulas in Γ_a and Γ are derivable, then one of the formulas in Δ_b and Δ is derivable. The policy context in the conclusion contains only w . However, formulas in Γ and Δ are wrapped under \Box_{γ} . Because submitting w and γ is the same as submitting w first, then submitting γ , if all of the formulas in $\Box_{\gamma} \Gamma$ are true if w is submitted, it implies that all of the formulas in Γ are true if w and γ is true. The same reasoning applies to Δ . Recall that positive formulas cannot be falsified by strengthening policies and negative formulas cannot be falsified by weakening policies. Formulas in Γ_a are positive, so if they are true under w , then they are true under w and γ . Formulas in Δ_b are negative, so if one of them is true under w and γ , then it is true under w as well. This is why the $\Box R$ rule is sound. The left rule removes the conditional γ if the policy γ is true already.

This formulation can be shown to be sound and complete with regard to the axiomatized system. However, cut elimination cannot be proven. This is similar to the problem of Prawitz's formulations

of necessity [35]. One way to solve this problem as pointed out by Pfenning et al. is to have separate judgments for truths and validity [34]. This is precisely what we do here. We have a separate judgment for counterfactual conditions. Instead of having a shared policy world, we index each formula with its policy world using judgment: $\varphi @ w$.

World Relations and Kripke Semantics Our proof rules explicitly index formulas with policy worlds and operate on policy worlds. There is a strong connection between the Kripke Semantics defined by Becker et al. [5] and how our proof rules manipulate policy worlds.

A Kripke model M is a triple $\langle W, R, V \rangle$, where W is a set of worlds, $R \subseteq \wp(W) \times W \times W$, and $V : \mathbf{At} \rightarrow \wp(W)$. We write $\wp(W)$ to denote the power set of W . The semantics of counterfactuals given the model is as follows:

$w \Vdash_M \Box_{\gamma} \varphi$ iff $\forall w', R|_{\gamma}|_M(w, w') \Rightarrow w' \Vdash_M \varphi$ where $|_{\gamma}|_M = \{w \in W \mid w \Vdash_M \gamma\}$.

Here, the accessibility relation between worlds are indexed by a set of worlds where a policy γ is true ($R|_{\gamma}|_M(w, w')$).

In general, a formula is valid if it is valid for all models. However, because the worlds are closely related to policies, not all models are reasonable models given how policies relate to each other. Becker et al. defined TM models, where each TM model models all possible policies and all possible policy interactions. One important interaction between policies are the relative strength of policies. This is what we use in our conversion rule as well. One interesting observation is that our $\Box R$ rule only picks one world that is accessible from w given γ , namely (w, γ) . The Kripke semantics require φ hold in all worlds related to w given γ ($R|_{\gamma}|_M(w, w')$). This is another reason why we need the conversion rule. Intuitively, this rule makes sure that checking φ true in one accessible world is good enough. This also suggests that we could modify our $\Box R$ to build in conversion so all accessible worlds are considered. We leave this for future work.

New judgement in the premise The POLCUT and CONV rules introduce a new judgement $p @ w$ into the premise. As a consequence, our sequent calculus does not have the traditional sub-formula property. However, because we only consider systems with a finite number of atomic propositions, there is only a finite number of such $p @ w$. Further, some of the proofs of the meta-properties (e.g., soundness) rely on the fact that the new judgment $p @ w$ is only introduced when it is derivable: in the case of POLCUT , $w \Vdash_d p$; and in the case of CONV , Γ already includes $p @ w'$ where w' is weaker than w . If we were to allow the general cut rule, those proofs would have been very difficult, if not impossible.

Applications to Other Domains In addition to security, Datalog has been widely used in many other domains. For instance, Datalog has been used to model networks and diagnose network configuration [14, 24]. The counterfactual reasoning framework could potentially be applied to these domains to diagnose and pinpoint specific configurations or nodes in the network such that if it weren't for those nodes, the error wouldn't have occurred.

Counterfactual reasoning is a meta-level reasoning about a system. Here, this system is Horn clauses. The derivability relation for Horn clauses is fairly simple, so it is a great starting point to build a

formal proof system for counterfactual reasoning. We are interested in generalizing this framework to reason about counterfactuals of other logic, for instance, a constructive authorization logic. We can still interpret $\Box_\gamma \varphi$ as by submitting additional rules γ , φ is derivable. Then the challenge is how to design proof rules to capture the counterfactual reasoning in the same way the DLOG rule does for Horn clauses. Another challenge is to identify the requirements of the worlds in the Kripke structure to make it compatible with the semantics of that logic.

7 RELATED WORK

Counterfactual reasoning has been studied by philosophers for decades [10, 22, 26] after Lewis's seminal book was published in 1973 [22]. Most closely related to our reading of the counterfactual conditions $\Box_\gamma \varphi$ is work on conditional logic [2, 3, 11, 17, 18, 21, 29, 30]. In conditional logic, $H > F$ is read as: if H were true then F would have been the case; or F follows from H in all worlds where H is true. $\Box_\gamma \varphi$ is similar to $\gamma > \varphi$. There is a set of axioms dictating how conditionals interact with the rest of logic. Each conditional logic admits a different set of axioms. Most of the axioms that our sequent calculus support can be mapped to axioms of conditional logic. One exception is the Dlog axiom, which is intrinsic to Horn clause reasoning. Conditional logic is given possible world semantics. Similar to the Kripke semantics that Becker et al. introduced, there is a notion of accessibility of worlds indexed by conditionals (H), and a notion of a minimal world or multiple minimal worlds that are accessible from w with distance H . Because we target Horn clauses, the accessibility relation of the Kripke semantics for our logic is customized to accommodate properties of Horn clauses. To some extent, our proof system can be viewed as a special instance of a more general conditional logic. The concrete semantics help us avoid many of the unsatisfactory explanations of axioms in a more general conditional logic.

Only relatively recently, applications of counterfactual and conditional reasoning have found their way into computer science. Samet [36] developed a reasoning system based on conditional logic for game theory. Players deduce their next moves and backtrack possible states by considering other possible worlds. Halpern later reformalized Samet's system by using a combination of epistemic and conditional logic [19]. Becker et al. used counterfactual reasoning to model probing attacks in trust management systems [5].

Our work is inspired by the axiomatized proof system developed by Becker et al. [5]. Unfortunately, axiomatized proof systems are known to be difficult for proof construction. To address this problem, most recently, Pasarella et al. used a logic programming language [27] as the operational framework for constructing proofs in counterfactual reasoning in trust management system [31]. Their extended semantics for logic programs, written $P \vdash_{\Delta} G$, model semantics of counterfactual conditions. The validity of a formula G satisfies the condition that it is not possible to find a policy Δ such that $\Delta \vdash_{\Delta} \neg G$. Different from our work, Pasarella and Lobo's do not provide sequent calculus formulation for counterfactual reasoning.

We believe we are the first to define a sequent calculus for counterfactual conditionals based on Horn clauses. Our sequent calculus has cut elimination. All proofs of our metatheories are clean and

easy to follow as they are either by induction over the structure of the formula, or by induction over the sequent calculus rules.

Our formalization is influenced by the formalization of classical modal S5 by Murphy et al. [28] and work on judgmental reconstruction of modal logic [34]. Here, the counterfactual conditional γ is treated as a world in the Kripke model and therefore, we can borrow ideas from sequent calculus for modal logic. Interestingly, our initial attempt that does not use a separate judgment $\varphi @ w$ failed in very similar ways as Prawitz's formulations of necessity [35].

8 CONCLUSION

We presented a sequent calculus for counterfactual reasoning of Horn clauses, which can model probing attacks in trust management systems concisely. Our sequent calculus has cut elimination [32] and is sound and complete with respect to the axiomatized system proposed by Becker et al. We plan to explore applications of our sequent calculus to model counterfactual reasonings in domains including security, networks, and program analysis.

ACKNOWLEDGEMENT

This research was supported in part by NSF grant CNS1423168.

REFERENCES

- [1] Serge Abiteboul, Richard Hull, and Victor Vianu. 1994. *Foundations of Databases*. Pearson, Cambridge, MA, USA.
- [2] Régis Alenda, Nicola Olivetti, and Gian Luca Pozzato. 2012. Nested Sequent Calculi for Conditional Logics. In *Logics in Artificial Intelligence (Lecture Notes in Computer Science)*, Vol. 7519. Springer, 14–27.
- [3] Horacio Arlo-Costa. 2014. The Logic of Conditionals. In *The Stanford Encyclopedia of Philosophy* (summer 2014 ed.), Edward N. Zalta (Ed.).
- [4] Moritz Becker, Cedric Fournet, and Andrew Gordon. 2007. Design and Semantics of a Decentralized Authorization Language. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF'07)*.
- [5] Moritz Y Becker, Alessandra Russo, and Nik Sultana. 2012. Foundations of logic-based trust management. In *Proceedings of the 2012 IEEE Symposium on Security and Privacy (SP '12)*.
- [6] Matt Blaze, Joan Feigenbaum, and Jack Lacy. 1996. Decentralized Trust Management. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy (SP '96)*.
- [7] Piero A. Bonatti. 2011. Datalog for Security, Privacy and Trust. In *Proceedings of the First International Conference on Datalog Reloaded (Datalog'10)*.
- [8] Martin Bravenboer and Yannis Smaragdakis. 2009. Strictly Declarative Specification of Sophisticated Points-to Analyses. In *Proceedings of the 24th ACM SIGPLAN Conference on Object Oriented Programming Systems Languages and Applications (OOPSLA '09)*.
- [9] Michael Burrows, Martin Abadi, and Roger Needham. 1990. A Logic of Authentication. *ACM Trans. Comput. Syst.* 8, 1 (Feb. 1990).
- [10] J. Collins, E. Hall, and L. Paul. 2004. *Causation and Counterfactuals*. MIT Press, Cambridge, Mass.
- [11] H C M de Swart. 1983. Gentzen- or Beth-Type System, a Practical Decision Procedure and a Constructive Completeness Proof for the Counterfactual Logics VC and VCS. *The Journal of Symbolic Logic* 48, 1 (March 1983), 1–20.
- [12] John DeTreville. 2002. Binder, a Logic-Based Security Language. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy (SP '02)*.
- [13] Henry DeYoung, Deepak Garg, and Frank Pfenning. 2008. An Authorization Logic With Explicit Time. In *Proceedings of the 2008 21st IEEE Computer Security Foundations Symposium (CSF '08)*.
- [14] Ari Fogel, Stanley Fung, Luis Pedrosa, Meg Walraed-Sullivan, Ramesh Govindan, Ratul Mahajan, and Todd Millstein. 2015. A General Approach to Network Configuration Analysis. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation (NSDI'15)*.
- [15] Deepak Garg and Martin Abadi. 2008. A Modal Deconstruction of Access Control Logics. In *Proceedings of the 11th International Conference on Foundations of Software Science and Computational Structures (FOSSACS'08)*.
- [16] Valerio Genovese, Laura Giordano, Valentina Gliozzi, and Gian Luca Pozzato. 2011. A Conditional Constructive Logic for Access Control and Its Sequent Calculus. In *Automated Reasoning with Analytic Tableaux and Related Methods (Lecture Notes in Computer Science)*, Vol. 6793. Springer, 164–179.

- [17] Ian Philip Gent. 1992. A Sequent- or Tableau-Style System for Lewis's Counterfactual Logic VC. *Notre Dame Journal of Formal Logic* 33, 3 (1992), 369–382.
- [18] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Camilla Schwind. 2009. Tableau Calculus for Preference-Based Conditional Logics: PCL and its Extensions. *ACM Transactions on Computational Logic* 10, 3 (April 2009), 1–50.
- [19] Joseph Y. Halpern. 2000. Hypothetical Knowledge and Counterfactual Reasoning. In *Very Large Data Bases*.
- [20] Trevor Jim. 2001. SD3: A Trust Management System with Certified Evaluation. In *Proceedings of the 2001 IEEE Symposium on Security and Privacy (SP'01)*.
- [21] Björn Lellmann and Dirk Pattinson. 2012. Sequent Systems for Lewis' Conditional Logics. In *Logics in Artificial Intelligence (Lecture Notes in Computer Science)*, Vol. 7519. Springer, 320–332.
- [22] David Lewis. 1973. *Counterfactuals*. Blackwell, Oxford.
- [23] Ninghui Li, Benjamin Grosf, and Joan Feigenbaum. 2000. A Practically Implementable and Tractable Delegation Logic. In *Proceedings of the 2000 IEEE Symposium on Security and Privacy (SP'00)*.
- [24] Boon Thau Loo, Tyson Condie, Mimos Garofalakis, David E. Gay, Joseph M. Hellerstein, Petros Maniatis, Raghu Ramakrishnan, Timothy Roscoe, and Ion Stoica. 2006. Declarative Networking: Language, Execution and Optimization. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data (SIGMOD '06)*.
- [25] McKenna McCall, Lay Kuan Loh, and Limin Jia. 2017. *A Sequent Calculus for Counterfactual Reasoning*. Technical Report CMU-CyLab-17-003. CyLab, Carnegie Mellon University.
- [26] Peter Menzies. 2014. Counterfactual Theories of Causation. In *The Stanford Encyclopedia of Philosophy* (spring 2014 ed.), Edward N. Zalta (Ed.).
- [27] Dale Miller. 1989. A Logical Analysis of Modules in Logic Programming. *The Journal of Logic Programming* 6, 1-2 (Jan. 1989), 79–108.
- [28] Tom Murphy, Karl Cray, and Robert Harper. 2005. Distributed Control Flow with Classical Modal Logic. In *Computer Science Logics (Lecture Notes in Computer Science)*, Vol. 3634. Springer, 51–69.
- [29] Nicola Olivetti and Gian Luca Pozzato. 2015. A Standard Internal Calculus for Lewis' Counterfactual Logics. In *Proceedings of the 24th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods - Volume 9323 (TABLEAUX 2015)*.
- [30] Nicola Olivetti, Gian Luca Pozzato, and Camilla B Schwind. 2007. A Sequent Calculus and a Theorem Prover for Standard Conditional Logics. *ACM Transactions on Computational Logic* 8, 4 (August 2007).
- [31] Edelmira Pasarella and Jorge Lobo. 2015. Reasoning about Policy Behavior in Logic-Based Trust Management Systems: Some Complexity Results and an Operational Framework. In *Proceedings of the 28th IEEE Computer Security Foundations Symposium (CSF'15)*.
- [32] Frank Pfenning. 1995. Structural Cut Elimination. In *Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science (LICS '95)*.
- [33] Frank Pfenning. 2000. Structural Cut Elimination: I. Intuitionistic and Classical Logic. *Information and Computation* 157, 1-2 (February 2000), 84–141.
- [34] Frank Pfenning and Rowan Davies. 2001. A Judgmental Reconstruction of Modal Logic. *Mathematical Structures in Computer Science* 11, 4 (Aug. 2001), 511–540.
- [35] D Prawitz. 1965. *Natural Deduction*. Almqvist Wiksell, Stockholm.
- [36] D. Samet. 1996. Hypothetical Knowledge and Games with Perfect Information. *Games and Economic Behavior* 17 (1996), 230–251.
- [37] Leon Sterling and Ehud Shapiro. 1986. *The Art of Prolog: Advanced Programming Techniques*. MIT Press, Cambridge, MA, USA.
- [38] Ankur Taly, Úlfar Erlingsson, John C. Mitchell, Mark S. Miller, and Jasvir Nagra. 2011. Automated Analysis of Security-Critical JavaScript APIs. In *Proceedings of the 2011 IEEE Symposium on Security and Privacy (SP '11)*.

A LEMMAS AND PROOFS FOR SOUNDNESS

Proving Lemma 7 is fairly involved. Inducting over the structure of $\mathcal{E} :: \gamma_1 @ \top \vdash_{\Sigma} \gamma_2 @ \top$ proves insufficient because POLCUT and CONV can introduce judgements that our induction hypothesis does not account for. We have to prove a more general lemma instead. First, we define a new calculus which contains labels that will allow us to track judgements which have been introduced by, or are derived from judgements introduced by, POLCUT and CONV . This will allow us to distinguish judgements of the form $\gamma @ \top$ from those of the form $p @ w$. We assign the label 'A' to judgements that have been added by, or are derived from judgements which have been added by, POLCUT^+ or CONV^+ . All other judgements are labeled with 'o'.

Our syntax is the same as in Section 4 with a few exceptions shown below.

Label	x	$:: =$	$A \mid o$
Judgement	$J, (J)_x$	$:: =$	$\gamma @ w \mid (\gamma @ w)_x$
Contexts	Γ^+	$:: =$	$\cdot \mid \Gamma^+, (J)_x$

The judgment for the marked up derivation is $\Gamma^+ \vdash_{\Sigma}^+ \Delta$. The sequent rules are shown in Figure 9. For the following lemmas, we define some notation:

$$\begin{aligned} \forall J \in \Gamma_p^+, J = (q @ w)_A & & \forall J \in \Gamma_p, J = q @ w \\ \forall J \in \Gamma_{\top}^+, J = (\gamma @ \top)_o & & \forall J \in (\Gamma_{\top}, \Delta_{\top}), J = \gamma @ \top \\ \forall J \in \Gamma_{\text{At}}^+, J = (p @ \vec{q})_o & & \forall J \in (\Gamma_{\text{At}}, \Delta_{\text{At}}), J = p @ \vec{q} \end{aligned}$$

We prove that derivations using original sequent rules can be annotated, which result in derivations in newly defined sequent rules.

Lemma 16. *If $\mathcal{E} :: \Gamma_p, \Gamma_{\top}, \Gamma_{\text{At}} \vdash_{\Sigma} \Delta_{\top}, \Delta_{\text{At}}$ then $\exists \mathcal{D} :: \Gamma_p^+, \Gamma_{\top}^+, \Gamma_{\text{At}}^+ \vdash_{\Sigma}^+ \Delta_{\top}, \Delta_{\text{At}}$*

Proof (sketch): By induction over the structure of \mathcal{E} . Most of the work for this proof involves showing that the induction hypothesis can be applied to the subderivations of \mathcal{E} . Once that has been shown, the conclusion follows from applying the same rule from the labeled version of the sequent calculus on the result(s) of applying the induction hypothesis. \square

Next, we relate derivations in the marked up sequent rules to formula semantics. We define $\text{polOf}(J)$ to transform a judgment to a policy as follows.

$$\begin{aligned} \text{polOf}(\gamma @ \top) &= \gamma \\ \text{polOf}(p @ \vec{q}) &= p \vdash \vec{q} \end{aligned}$$

Lemma 17. *If $\mathcal{E} :: \Gamma_p^+, \Gamma_{\top}^+, \Gamma_{\text{At}}^+ \vdash_{\Sigma}^+ \Delta_{\top}, \Delta_{\text{At}}$ where $\forall J \in \Gamma_p^+, J = q @ w$ and $\text{polOf}(\Gamma_{\top}^+, \Gamma_{\text{At}}^+, w) \Vdash q$, then $\exists J \in (\Delta_{\top}, \Delta_{\text{At}})$ s.t. $\text{polOf}(\Gamma_{\top}^+, \Gamma_{\text{At}}^+) \Vdash \text{polOf}(J)$*

Proof (sketch): By induction over the structure of \mathcal{E} . Most of the work for this proof involves showing that the induction hypothesis can be applied to the subderivations of \mathcal{E} . Then, the conclusion follows directly, or from one of the lemmas from Appendix B. \square

Lemma 7. *$\forall \gamma_1, \gamma_2$, If $\mathcal{E} :: \gamma_1 @ \top \vdash_{\Sigma} \gamma_2 @ \top$ then $\bar{\gamma}_1 \Vdash \gamma_2$.*

Proof (sketch): The proof follows from the observation that \mathcal{E} is of the form $\Gamma_{\top} \vdash_{\Sigma} \Delta_{\top}$. Then the conclusion follows from, Lemma 16 and Lemma 17. \square

B PROPERTIES OF \Vdash AND \Vdash_d

Lemma 18 (Admissibility of Cut for \Vdash). *If $\mathcal{D} :: w \Vdash \gamma_1$ and $\mathcal{E} :: w, \bar{\gamma}_1 \Vdash \gamma_2$ then $\exists \mathcal{F} :: w \Vdash \gamma_2$*

Proof (sketch): If we consider the definitions for \Vdash as sequent, this is a proof of admissibility of cut. So we proceed by induction first over the structure of γ_1 and then over the structures of \mathcal{D} and \mathcal{E} . The proof closely resembles our proof of Theorem 4 and requires weakening, contraction for \Vdash . Most cases are straightforward, except for the one where \mathcal{D} ends in $\square R$ and \mathcal{E} ends in At which requires Lemma 23. We also need Lemma 22 to prove the admissibility of cut for \Vdash_d . \square

$$\begin{array}{c}
\frac{}{\Gamma^+, (p @ w)_x \vdash_{\Sigma}^+ p @ w, \Delta} \text{INIT}^+ \quad \frac{}{\Gamma^+ \vdash_{\Sigma}^+ \top @ w, \Delta} \top \text{R}^+ \quad \frac{\Gamma^+ \vdash_{\Sigma}^+ \gamma_1 @ w, \gamma_1 \wedge \gamma_2 @ w, \Delta \quad \Gamma^+ \vdash_{\Sigma}^+ \gamma_2 @ w, \gamma_1 \wedge \gamma_2 @ w, \Delta}{\Gamma^+ \vdash_{\Sigma}^+ \gamma_1 \wedge \gamma_2 @ w, \Delta} \wedge \text{R}^+ \\
\frac{\Gamma^+, (\gamma_1 @ w)_o, (\gamma_1 \wedge \gamma_2 @ w)_o, \vdash_{\Sigma}^+ \Delta}{\Gamma^+, (\gamma_1 \wedge \gamma_2 @ w)_o \vdash_{\Sigma}^+ \Delta} \wedge \text{L1}^+ \quad \frac{\Gamma^+, (\gamma_2 @ w)_o, (\gamma_1 \wedge \gamma_2 @ w)_o, \vdash_{\Sigma}^+ \Delta}{\Gamma^+, (\gamma_1 \wedge \gamma_2 @ w)_o \vdash_{\Sigma}^+ \Delta} \wedge \text{L2}^+ \quad \frac{\Gamma^+ \vdash_{\Sigma}^+ \varphi @ (w, \bar{\gamma}), \Box_{\gamma} \varphi @ w, \Delta}{\Gamma^+ \vdash_{\Sigma}^+ \Box_{\gamma} \varphi @ w, \Delta} \Box \text{R}^+ \\
\frac{\Gamma^+, (\Box_{\gamma} \varphi @ w)_x, (\varphi @ (w, \bar{\gamma}))_x \vdash_{\Sigma}^+ \Delta}{\Gamma^+, (\Box_{\gamma} \varphi @ w)_x \vdash_{\Sigma}^+ \Delta} \Box \text{L}^+ \quad \frac{w \Vdash_d p \quad \text{symbols}(p @ w) \subseteq \Sigma \quad \Gamma^+, (p @ w)_A \vdash_{\Sigma}^+ \Delta}{\Gamma^+ \vdash_{\Sigma}^+ \Delta} \text{POLCUT}^+ \\
\frac{\Gamma^+, (p @ (w, w', q))_x, (q @ w)_y, (p @ (w, w'))_x \vdash_{\Sigma}^+ \Delta}{\Gamma^+, (p @ (w, w', q))_x, (q @ w)_y \vdash_{\Sigma}^+ \Delta} \text{POLL-W-AT}^+ \\
\frac{\Gamma^+, (p @ (w, w', \Box_{\vec{a}} b))_A, (b @ (w, \vec{a}))_x, (p @ (w, w'))_A \vdash_{\Sigma}^+ \Delta}{\Gamma^+, (p @ (w, w', \Box_{\vec{a}} b))_A, (b @ (w, \vec{a}))_x \vdash_{\Sigma}^+ \Delta} \text{POLL-W-}\Box^+ \\
\frac{\Gamma^+, (s @ (q :- p_1, \dots, p_n, w))_A, (s @ w)_A \vdash_{\Sigma}^+ \Delta \quad \Gamma^+, (s @ (q :- p_1, \dots, p_n, w))_A, (s @ (q, w))_A, (p_1 @ w)_A, \dots, (p_n @ w)_A \vdash_{\Sigma}^+ \Delta}{\Gamma^+, (s @ (q :- p_1, \dots, p_n, w))_A \vdash_{\Sigma}^+ \Delta} \text{DLOG}^+ \\
\frac{\text{symbols}(w_2) \subseteq \Sigma \quad \Gamma^+, (p @ w_1)_x, (p @ w_2)_A \vdash_{\Sigma}^+ \Delta \quad (\wedge(w_2) @ \top)_o \vdash_{\Sigma}^+ \wedge(w_1) @ \top}{\Gamma^+, (p @ w_1)_x \vdash_{\Sigma}^+ \Delta} \text{CONV}^+
\end{array}$$

Figure 9: Labeled sequent calculus

Lemma 19. $\forall w, \text{If } \mathcal{E} :: w, p :- q_1 \cdots q_n \Vdash_d s \text{ then}$

- (1) $\exists \mathcal{D} :: w \Vdash_d s$, or
- (2) $\exists \mathcal{D} :: w, p \Vdash_d s$, and $\forall i \in [1, n] \exists \mathcal{D}_i :: w \Vdash_d q_i$

Proof (sketch): By induction over the structure of \mathcal{E} . When \mathcal{E} ends in AT, the conclusion follows directly from the AT rule.

When \mathcal{E} ends in DLOG and $p = s$, the conclusion follows from AT and the induction hypothesis. Otherwise, when $p \neq s$, the conclusion follows from the induction hypothesis and DLOG. \square

Lemma 20 (Contraction for \Vdash). $\forall w, w', \gamma \text{ If } \mathcal{E} :: w, w', w' \Vdash \gamma \text{ then } \exists \mathcal{D} :: w, w' \Vdash \gamma$

Proof (sketch): By induction over the structure of \mathcal{E} . We can get the conclusion by applying the same rule on the result of using the induction hypothesis. Lemma 24 is required to prove contraction for \Vdash_d . \square

Lemma 21 (Weakening for \Vdash). $\forall w, w', \gamma \text{ If } \mathcal{E} :: w' \Vdash \gamma \text{ then } \exists \mathcal{D} :: w, w' \Vdash \gamma \text{ with } |\mathcal{D}| = |\mathcal{E}|$

Proof (sketch): By induction over the structure of \mathcal{E} . We can get the conclusion by applying the same rule on the result of using the induction hypothesis. Lemma 25 is required to prove weakening for \Vdash_d . \square

Lemma 22 (Admissibility of Cut for \Vdash_d). $\forall w \text{ If } \mathcal{D} :: w \Vdash_d p \text{ and } \mathcal{E} :: w, p \Vdash_d q \text{ then } \exists \mathcal{F} :: w \Vdash_d q$

Proof (sketch): By nested induction on the structures of \mathcal{D} and \mathcal{E} . The conclusion follows from contraction when \mathcal{D} ends in AT and

\mathcal{E} is arbitrary and from the AT rule when \mathcal{D} is arbitrary and \mathcal{E} ends in AT. The last case, when \mathcal{D} and \mathcal{E} both end in DLOG, follows from applying DLOG on the result of applying the induction hypothesis on \mathcal{D} and the subderivations of \mathcal{E} . \square

Lemma 23. $\forall w \text{ If } \mathcal{D} :: w, \vec{q} \Vdash_d r \text{ and } \mathcal{E} :: w, \Box_{\vec{q}} r \Vdash_d p \text{ then } \exists \mathcal{F} :: w \Vdash_d p$

Proof (sketch): By nested induction over the structures of \mathcal{D} and \mathcal{E} . The conclusion follows from the AT rule when \mathcal{D} is arbitrary and \mathcal{E} ends in AT. When \mathcal{D} is arbitrary and \mathcal{E} ends in DLOG, the conclusion follows from applying DLOG on the results of applying the induction hypothesis on \mathcal{D} and the subderivations of \mathcal{E} . \square

Lemma 24 (Contraction for \Vdash_d). $\forall w, w', \gamma \text{ If } \mathcal{E} :: w, w', w' \Vdash_d \gamma \text{ then } \exists \mathcal{D} :: w, w' \Vdash_d \gamma$

Proof (sketch): By induction over the structure of \mathcal{E} . When \mathcal{E} ends in AT, the conclusion follows directly from the AT rule. When \mathcal{E} ends in DLOG, we can get the conclusion by applying the DLOG on the result of using the induction hypothesis on the subderivations of \mathcal{E} . \square

Lemma 25 (Weakening for \Vdash_d). $\forall w, w', \gamma \text{ If } \mathcal{E} :: w' \Vdash_d \gamma \text{ then } \exists \mathcal{D} :: w, w' \Vdash_d \gamma \text{ with } |\mathcal{D}| = |\mathcal{E}|$

Proof (sketch): By induction over the structure of \mathcal{E} . When \mathcal{E} ends in AT, the conclusion follows directly from the AT rule. When \mathcal{E} ends in DLOG, we can get the conclusion by applying the DLOG on the result of using the induction hypothesis on the subderivations of \mathcal{E} . \square

C ADMISSIBILITY OF CUT

Theorem 4 (Admissibility of Cut).

If $\mathcal{D} :: \Gamma \vdash \varphi @ w, \Delta$ and $\mathcal{E} :: \Gamma, \varphi @ w \vdash \Delta$,
then $\exists \mathcal{F} :: \Gamma \vdash \Delta$

PROOF.

We show some representative cases below. The ones not shown proceed in the same fashion.

Base cases

Case: \mathcal{D} ends in INIT and \mathcal{E} is arbitrary.

$$\mathcal{D} = \frac{}{\Gamma, p @ w \vdash_\Sigma p @ w, \Delta} \text{INIT}$$

and $\mathcal{E} :: \Gamma, p @ w, p @ w \vdash_\Sigma \Delta$

To show: $\exists \mathcal{F} :: \Gamma, p @ w \vdash_\Sigma \Delta$

\mathcal{F} may be obtained by applying the contraction lemma on \mathcal{E} .

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in INIT.

$$\mathcal{D} :: \Gamma \vdash_\Sigma p @ w, p @ w, \Delta$$

$$\text{and } \mathcal{E} = \frac{}{\Gamma, p @ w \vdash_\Sigma p @ w, \Delta} \text{INIT}$$

To show: $\exists \mathcal{F} :: \Gamma \vdash_\Sigma p @ w, \Delta$

\mathcal{F} may be obtained by applying the contraction lemma on \mathcal{D} .

Case: \mathcal{D} ends in \top R and \mathcal{E} is arbitrary.

$\top @ w$ will be a side formula in \mathcal{E} .

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in \top R.

The cut formula is a side formula in \mathcal{E} .

Case: \mathcal{D} ends in \perp L and \mathcal{E} is arbitrary.

The cut formula is a side formula in \mathcal{D} .

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in \perp L.

$\perp @ w$ is a side formula in \mathcal{D}

Principal cases

Case: \mathcal{D} ends in \Box R and \mathcal{E} ends in \Box L

$$\mathcal{D}_1 :: \Gamma \vdash_\Sigma \varphi @ (w, \bar{\gamma}), \Box_\gamma \varphi @ w, \Delta$$

$$\mathcal{D} = \frac{}{\Gamma \vdash_\Sigma \Box_\gamma \varphi @ w, \Delta} \Box R$$

$$\mathcal{E}_1 :: \Gamma, \Box_\gamma \varphi @ w, \varphi @ (w, \bar{\gamma}) \vdash_\Sigma \Delta$$

$$\mathcal{E} = \frac{}{\Gamma, \Box_\gamma \varphi @ w \vdash_\Sigma \Delta} \Box L$$

(G) To show: $\exists \mathcal{F} :: \Gamma \vdash_\Sigma \Delta$

Weakening and IH on $\Box_\gamma \varphi @ w, \mathcal{D}$, and \mathcal{E}_1 gives

(1) $\mathcal{E}'_1 :: \Gamma, \varphi @ (w, \bar{\gamma}) \vdash_\Sigma \Delta$

Weakening and IH on $\Box_\gamma \varphi @ w, \mathcal{D}_1$, and \mathcal{E} gives

(2) $\mathcal{D}'_1 :: \Gamma \vdash_\Sigma \varphi @ (w, \bar{\gamma}), \Delta$

IH on $\varphi @ (w, \bar{\gamma}), \mathcal{D}'_1$, and \mathcal{E}'_1 gives (G)

Special cases

These cases are already covered by the base cases and side cases.

We show one case in more detail, here, to make this explicit.

Case: \mathcal{E} ends in CONV

Subcase: \mathcal{D} ends in INIT

$$\mathcal{D} = \frac{}{\Gamma, p @ w_1 \vdash_\Sigma p @ w_1, \Delta} \text{INIT}$$

$$\mathcal{E}_1 :: \text{symbols}(w_2) \subseteq \Sigma$$

$$\mathcal{E}_2 :: \Gamma, p @ w_1, p @ w_1, p @ w_2 \vdash_\Sigma \Delta$$

$$\mathcal{E}_3 :: \wedge(w_2) @ \top \vdash_\Sigma \wedge(w_1) @ \top$$

$$\mathcal{E} = \frac{}{\Gamma, p @ w_1, p @ w_1 \vdash_\Sigma \Delta} \text{CONV}$$

To show: $\exists \mathcal{F} :: \Gamma, p @ w_1 \vdash_\Sigma \Delta$

Applying the contraction lemma on \mathcal{E} gives \mathcal{F}

The rest of the subcases are similar. We show one below.

Subcase: \mathcal{D} ends in \vee R1

Observe that $p @ w_1$ must be a side formula in \mathcal{D} .

$$\mathcal{D}_1 :: \Gamma \vdash_\Sigma p @ w_1, \varphi_1 @ w, \varphi_1 \vee \varphi_2 @ w, \Delta$$

$$\mathcal{D} = \frac{}{\Gamma \vdash_\Sigma p @ w_1, \varphi_1 \vee \varphi_2 @ w, \Delta} \vee R1$$

$$\mathcal{E}_1 :: \text{symbols}(w_2) \subseteq \Sigma$$

$$\mathcal{E}_2 :: \Gamma, p @ w_1, p @ w_2 \vdash_\Sigma \varphi_1 \vee \varphi_2 @ w, \Delta$$

$$\mathcal{E}_3 :: \wedge(w_2) @ \top \vdash_\Sigma \wedge(w_1) @ \top$$

$$\mathcal{E} = \frac{}{\Gamma, p @ w_1 \vdash_\Sigma \varphi_1 \vee \varphi_2 @ w, \Delta} \text{CONV}$$

(G) To show: $\exists \mathcal{F} :: \Gamma \vdash_\Sigma \varphi_1 \vee \varphi_2 @ w, \Delta$

Apply weakening and IH on $p @ w_1, \mathcal{D}_1$, and \mathcal{E} gives

(1) $\mathcal{D}'_1 :: \Gamma \vdash_\Sigma \varphi_1 @ w, \varphi_1 \vee \varphi_2 @ w, \Delta$

Applying $\vee R1$ on \mathcal{D}'_1 gives (G)

Side cases on \mathcal{D}

These cases are all similar to each other.

We show all of the cases for our new rules, except $\Box R$, $\Box L$, and $\text{POL-L-W-}\Box$ plus one involving INIT.

Case: \mathcal{D} ends in INIT and \mathcal{E} is arbitrary

$$\mathcal{D} = \frac{}{\Gamma, p @ w \vdash_\Sigma \varphi' @ w', p @ w, \Delta} \text{INIT}$$

and $\mathcal{E} :: \Gamma, p @ w, \varphi' @ w' \vdash_\Sigma p @ w, \Delta$

To show: $\exists \mathcal{F} :: \Gamma, p @ w \vdash_\Sigma p @ w, \Delta$

Let $\mathcal{F} = \frac{}{\Gamma, p @ w \vdash_\Sigma p @ w, \Delta} \text{INIT}$

Case: \mathcal{D} ends in POLCUT and \mathcal{E} is arbitrary

$$\mathcal{D}_1 :: w \Vdash_d p \quad \mathcal{D}_2 :: \text{symbols}(p @ w) \subseteq \Sigma$$

$$\mathcal{D}_3 :: \Gamma, p @ w \vdash_\Sigma \varphi' @ w', \Delta$$

$$\mathcal{D} = \frac{}{\Gamma \vdash_\Sigma \varphi' @ w', \Delta} \text{POLCUT}$$

and $\mathcal{E} :: \Gamma, \varphi' @ w' \vdash_\Sigma \Delta$

(G) To show: $\exists \mathcal{F} :: \Gamma \vdash_\Sigma \Delta$

Weakening and IH on $\varphi' @ w', \mathcal{D}_3$, and \mathcal{E} gives

(1) $\mathcal{D}'_3 :: \Gamma, p @ w \vdash_\Sigma \Delta$

Applying POLCUT on $\mathcal{D}_1, \mathcal{D}_2$, and \mathcal{D}'_3 gives (G)

Case: \mathcal{D} ends in POL-L-W-AT and \mathcal{E} is arbitrary

$$\frac{\mathcal{D}_1 :: \Gamma, p @ (w, w', q), q @ w, \\ p @ (w, w') \vdash_{\Sigma} \varphi' @ w_1, \Delta}{\mathcal{D} = \Gamma, p @ (w, w', q), q @ w \vdash_{\Sigma} \varphi' @ w_1, \Delta} \text{PoLL-W-AT}$$

and $\mathcal{E} :: \Gamma, p @ (w, w', q), q @ w, \varphi' @ w_1 \vdash_{\Sigma} \Delta$

(G) To show: $\exists \mathcal{F} :: \Gamma, p @ (w, w', q), q @ w \vdash_{\Sigma} \Delta$
Weakening and IH on $\varphi' @ w_1, \mathcal{D}_1$, and \mathcal{E} gives

$$(1) \mathcal{D}'_1 :: \Gamma, p @ (w, w', q), q @ w, p @ (w, w') \vdash_{\Sigma} \Delta$$

Applying PoLL-W-AT on \mathcal{D}'_1 gives (G)

Case: \mathcal{D} ends in DLOG and \mathcal{E} is arbitrary

$$\frac{\mathcal{D}_1 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ w \vdash_{\Sigma} \varphi' @ w', \Delta \\ \mathcal{D}_2 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ (q, w), \\ p_1 @ w, \dots, p_n @ w \vdash_{\Sigma} \varphi' @ w', \Delta}{\mathcal{D} = \Gamma, s @ (q :- p_1, \dots, p_n, w) \vdash_{\Sigma} \varphi' @ w', \Delta} \text{DLOG}$$

and $\mathcal{E} :: \Gamma, s @ (q :- p_1, \dots, p_n, w), \varphi' @ w' \vdash_{\Sigma} \Delta$

(G) To show: $\exists \mathcal{F} :: \Gamma, s @ (q :- p_1, \dots, p_n, w) \vdash_{\Sigma} \Delta$
Weakening and IH on $\varphi' @ w', \mathcal{D}_1$, and \mathcal{E} gives

$$(1) \mathcal{D}'_1 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ w \vdash_{\Sigma} \Delta$$

Weakening and IH on $\varphi' @ w', \mathcal{D}_2$, and \mathcal{E} gives

$$(2) \mathcal{D}'_2 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ (q, w), \\ p_1 @ w, \dots, p_n @ w \vdash_{\Sigma} \Delta$$

Applying DLOG on \mathcal{D}'_1 and \mathcal{D}'_2 gives (G)

Case: \mathcal{D} ends in CONV and \mathcal{E} is arbitrary

$$\frac{\mathcal{D}_1 :: \text{symbols}(w_2) \subseteq \Sigma \\ \mathcal{D}_2 :: \Gamma, p @ w_1, p @ w_2 \vdash_{\Sigma} \varphi' @ w, \Delta \\ \mathcal{D}_3 :: \wedge(w_2) @ \top \vdash_{\Sigma} \wedge(w_1) @ \top}{\mathcal{D} = \Gamma, p @ w_1 \vdash_{\Sigma} \varphi' @ w, \Delta} \text{CONV}$$

and $\mathcal{E} :: \Gamma, p @ w_1, \varphi' @ w \vdash_{\Sigma} \Delta$

(G) To show: $\exists \mathcal{F} :: \Gamma, p @ w_1 \vdash_{\Sigma} \Delta$
Weakening and IH on $\varphi' @ w, \mathcal{D}_2$, and \mathcal{E} gives

$$(1) \mathcal{D}'_2 :: \Gamma, p @ w_1, p @ w_2 \vdash_{\Sigma} \Delta$$

Applying CONV on $\mathcal{D}_1, \mathcal{D}'_2$, and \mathcal{D}_3 gives (G)

Side cases on \mathcal{E}

These cases are all similar to each other.

We show all of the cases for our new rules, except $\square R$,

$\square L$, and PoLL-W- \square , plus one involving INIT

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in INIT

$$\frac{\mathcal{D} :: \Gamma, p @ w \vdash_{\Sigma} \varphi' @ w', p @ w, \Delta \text{ and} \\ \mathcal{E} = \Gamma, p @ w, \varphi' @ w' \vdash_{\Sigma} p @ w, \Delta}{\text{INIT}}$$

To show: $\exists \mathcal{F} :: \Gamma, p @ w \vdash_{\Sigma} p @ w, \Delta$

Let $\mathcal{F} = \Gamma, p @ w \vdash_{\Sigma} p @ w, \Delta$

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in PoLCUT

$$\frac{\mathcal{D} :: \Gamma \vdash_{\Sigma} \varphi' @ w, \Delta \text{ and} \\ \mathcal{E}_1 :: w \Vdash_d p \quad \mathcal{E}_2 :: \text{symbols}(p @ w) \subseteq \Sigma \\ \mathcal{E}_3 :: \Gamma, p @ w, \varphi' @ w' \vdash_{\Sigma} \Delta}{\mathcal{E} = \Gamma, \varphi' @ w' \vdash_{\Sigma} \Delta} \text{PoLCUT}$$

(G) To show: $\exists \mathcal{F} :: \Gamma \vdash_{\Sigma} \Delta$

Weakening and IH on $\varphi' @ w', \mathcal{D}$, and \mathcal{E}_3 gives

$$(1) \mathcal{E}'_3 :: \Gamma, p @ w \vdash_{\Sigma} \Delta$$

Applying PoLCUT on $\mathcal{E}_1, \mathcal{E}_2$, and \mathcal{E}'_3 gives (G)

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in PoLL-W-AT

$$\mathcal{D} :: \Gamma, p @ (w, w', q), q @ w \vdash_{\Sigma} \varphi' @ w_1, \Delta \text{ and}$$

$$\mathcal{E}_1 :: \Gamma, p @ (w, w', q), q @ w, p @ (w, w'),$$

$$\varphi' @ w_1 \vdash_{\Sigma} \Delta$$

$$\mathcal{E} = \frac{\Gamma, p @ (w, w', q), q @ w, \varphi' @ w_1 \vdash_{\Sigma} \Delta}{\text{PoLL-W-AT}}$$

(G) To show: $\exists \mathcal{F} :: \Gamma, p @ (w, w', q), q @ w \vdash_{\Sigma} \Delta$

Weakening and IH on $\varphi' @ w_1, \mathcal{D}$, and \mathcal{E}_1 gives

$$(1) \mathcal{E}'_1 :: \Gamma, p @ (w, w', q), q @ w, p @ (w, w') \vdash_{\Sigma} \Delta$$

Applying PoLL-W-AT on \mathcal{E}'_1 gives (G)

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in DLOG

$$\mathcal{D} :: \Gamma, s @ (q :- p_1, \dots, p_n, w) \vdash_{\Sigma} \varphi' @ w', \Delta \text{ and}$$

$$\mathcal{E}_1 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ w, \varphi' @ w' \vdash_{\Sigma} \Delta$$

$$\mathcal{E}_2 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ (q, w),$$

$$p_1 @ w, \dots, p_n @ w, \varphi' @ w' \vdash_{\Sigma} \Delta$$

$$\mathcal{E} = \frac{\Gamma, s @ (q :- p_1, \dots, p_n, w), \varphi' @ w' \vdash_{\Sigma} \Delta}{\text{DLOG}}$$

(G) To show: $\exists \mathcal{F} :: \Gamma, s @ (q :- p_1, \dots, p_n, w) \vdash_{\Sigma} \Delta$

Weakening and IH on $\varphi' @ w', \mathcal{D}$, and \mathcal{E}_1 gives

$$(1) \mathcal{E}'_1 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ w \vdash_{\Sigma} \Delta$$

Weakening and IH on $\varphi' @ w', \mathcal{D}$, and \mathcal{E}_2 gives

$$(2) \mathcal{E}'_2 :: \Gamma, s @ (q :- p_1, \dots, p_n, w), s @ (q, w), \\ p_1 @ w, \dots, p_n @ w \vdash_{\Sigma} \Delta$$

Applying DLOG on \mathcal{E}'_1 and \mathcal{E}'_2 gives (G)

Case: \mathcal{D} is arbitrary and \mathcal{E} ends in CONV

$$\mathcal{D} :: \Gamma, p @ w_1 \vdash_{\Sigma} \varphi' @ w, \Delta \text{ and}$$

$$\mathcal{E}_1 :: \text{symbols}(w_2) \subseteq \Sigma$$

$$\mathcal{E}_2 :: \Gamma, p @ w_1, p @ w_2, \varphi' @ w \vdash_{\Sigma} \Delta$$

$$\mathcal{E}_3 :: \wedge(w_2) @ \top \vdash_{\Sigma} \wedge(w_1) @ \top$$

$$\mathcal{E} = \frac{\Gamma, p @ w_1, \varphi' @ w \vdash_{\Sigma} \Delta}{\text{CONV}}$$

(G) To show: $\exists \mathcal{F} :: \Gamma, p @ w_1 \vdash_{\Sigma} \Delta$

Weakening and IH on $\varphi' @ w, \mathcal{D}$, and \mathcal{E}_2 gives

$$(1) \mathcal{E}'_2 :: \Gamma, p @ w_1, p @ w_2 \vdash_{\Sigma} \Delta$$

Applying CONV on $\mathcal{E}_1, \mathcal{E}'_2$, and \mathcal{E}_3 gives (G) \square

D PROOFS FOR CONSISTENCY

Theorem 5 (Consistency).

$\forall n, \# \mathcal{E}, \# w$ such that $\mathcal{E} :: \Gamma \vdash \perp @ w$ and $|\mathcal{E}| \leq n$ where $\forall J \in \Gamma, J = p @ w$.

Proof (sketch): By induction on n . When $n = 0$, the conclusion holds by inversion. Assume the conclusion holds when $n < k$. Consider the case when $n = k$. The only possible derivation is one that ends in PoLCUT, PoLL-W-AT, PoLL-W- \square , DLOG, or CONV because of the restrictions on the judgements in Γ . However, that would yield a proof of $\mathcal{E}' :: \Gamma \vdash \perp @ w'$ and $|\mathcal{E}'| = k - 1$ which contradicts with our assumption that $\# \mathcal{E}, \# w$ such that $\mathcal{E} :: \Gamma \vdash \perp @ w$ and $|\mathcal{E}| < k$. \square

E LEMMAS FOR COMPLETENESS

Lemma 10 (Axiom C1). $\forall \gamma, w, \vdash_{\Sigma} \gamma @ (\bar{v}, w)$.

Proof (sketch): By induction over the structure of γ . When γ is of the form \top , the conclusion follows from the TR rule. For the cases where γ is of the form $p \in \text{At}$ or $q : -\bar{p}$, the conclusion follows

from the POLCUT rule. Finally, when γ is of the form $\gamma_1 \wedge \gamma_2$, the conclusion follows from the use of the induction hypothesis on γ_1 and γ_2 and applying $\wedge R$ on the weakened result. \square

Lemma 26 (Generalized Axiom C2). *For all φ, γ, w, w' ,*

- (1) $\varphi @ (w, w', \bar{\gamma}), \gamma @ w \vdash_{\Sigma} \varphi @ (w, w')$ and
- (2) $\varphi @ (w, w'), \gamma @ w \vdash_{\Sigma} \varphi @ (w, w', \bar{\gamma})$.

Proof (sketch): By induction over the structure of φ . We prove (1) and (2) simultaneously. When φ is of the form \top or \perp , the conclusion follows directly from the $\top R$ and $\perp L$ rules, respectively. When φ is of the form $p \in \text{At}$, (1) follows from the POLL-W-AT rule and (2) from the CONV rule. The proofs for the rest of the cases where φ involves some connective follow a similar pattern. The conclusion follows from applying the left and right rules for the connective on the result(s) of applying the induction hypothesis on subformulas of φ . \square

Lemma 27 (Generalized Axiom Dlog).

$\forall w, \varphi, p, \vec{q}, \varphi$ is \square free, then

- (1) $\varphi @ (w, q :- \vec{p}), \vec{p} \rightarrow q @ w \vdash_{\Sigma} \varphi @ w$
- (2) $\varphi @ w, \vec{p} \rightarrow q @ w \vdash_{\Sigma} \varphi @ (w, q :- \vec{p})$

Proof (sketch): By induction over the structure of φ . We prove (1) and (2) simultaneously. When φ is of the form \top or \perp , the conclusion follows directly from the $\top R$ and $\perp L$ rules, respectively. When φ is of the form $S \in \text{At}$, (1) follows from the DLOG and $\rightarrow R$ rules, and (2) follows from CONV . The proofs for the rest of the cases where φ involves some connective follow a similar pattern. The conclusion follows from applying the left and right rules for the connective on the result(s) of applying the induction hypothesis on subformulas of φ . \square

Lemma 13 (Axiom MP). *If $\mathcal{E}_1 :: \cdot \vdash \varphi_1 @ w$ and $\mathcal{E}_2 :: \cdot \vdash \varphi_1 \rightarrow \varphi_2 @ w$ then $\exists \mathcal{E} :: \cdot \vdash \varphi_2 @ w$*

Proof (sketch): $\mathcal{D} :: \varphi_1 @ w, \varphi_1 \rightarrow \varphi_2 @ w \vdash_{\Sigma} \varphi_2 @ w$ may be constructed by applying $\rightarrow L$ and weakening on derivations from INIT . The conclusion follows by cutting $\varphi_1 @ w$ and $\varphi_1 \rightarrow \varphi_2 @ w$ from \mathcal{D} . \square

Lemma 28 (Generalized Axiom N).

If $\mathcal{E} :: \Gamma \vdash_{\Sigma} \Delta$, then $\exists \mathcal{D} :: \Gamma @ w \vdash_{\Sigma} \Delta @ w$

Where we define $\Gamma @ w$ as follows.

$$\begin{aligned} \cdot @ w &= \cdot \\ (\Gamma, \varphi @ w') @ w &= \Gamma @ w, \varphi @ (w, w') \end{aligned}$$

Proof (sketch): By induction over the structure of \mathcal{E} . For the cases where \mathcal{E} ends in INIT , $\top R$, or $\perp L$, the conclusion follows directly

from the INIT , $\top R$, and $\perp L$ rules, respectively. Most of the rest of the cases proceed the same way: the conclusion follows from applying the same rule on the result of applying the induction hypothesis on the subderivations of \mathcal{E} . The only exception is when \mathcal{E} ends in CONV , where an additional derivation comes from applying $\wedge L1$, $\wedge L2$, and weakening on a derivation from INIT . \square

The proof of Axiom Mon follows from a more general lemma, which requires a few additional lemmas to prove.

Lemma 15 (Axiom Mon.) *If $\vdash_{\Sigma} \gamma_1 \rightarrow \gamma_2 @ \top$ then $\forall w$*

$$\vdash_{\Sigma} \square_{\gamma_2} \varphi^+ \rightarrow \square_{\gamma_1} \varphi^+ @ w.$$

Proof (sketch): $\mathcal{D} :: \gamma_1 \wedge (\wedge w) @ \top \vdash (\gamma_2 \wedge (\wedge w)) @ \top$ may be obtained by applying $\wedge R$ on derivations obtained from Lemma 30 and application of INIT , weakening, and $\wedge R$. The conclusion follows from applying weakening, $\square R$, $\square L$, and $\rightarrow R$ on the result of applying Lemma 31 on \mathcal{D} . \square

Lemma 29.

If $\mathcal{E} :: \Gamma, \gamma_1 @ w \vdash_{\Sigma} \gamma_2 @ w, \Delta, \Delta_{\rightarrow}$, where $\forall J \in \Delta_{\rightarrow}, J = \gamma_1 \rightarrow \gamma_2 @ w$, then $\forall w, \Gamma, \gamma_1 @ w \vdash_{\Sigma} \gamma_2 @ w, \Delta$.

Proof (sketch): By induction over the structure of \mathcal{E} . \square

Lemma 30. *If $\mathcal{E} :: \vdash \gamma_1 \rightarrow \gamma_2 @ \top$ then $\gamma_1 @ \top \vdash \gamma_2 @ \top$.*

Proof (sketch): By weakening, we have $\gamma_1 @ \top \vdash \gamma_2 @ \top, \gamma_1 \rightarrow \gamma_2 @ \top$. Then, we can apply Lemma 29, with $\Gamma = \Delta = \cdot$ to reach the conclusion. \square

Lemma 31 (Generalized Axiom Mon).

If $\gamma_1 @ \top \vdash_{\Sigma} \gamma_2 @ \top$ then $\forall w$,

- (1) $\varphi^+ @ (w, \gamma_2) \vdash_{\Sigma} \varphi^+ @ (w, \gamma_1)$ and
- (2) $\varphi^- @ (w, \gamma_1) \vdash_{\Sigma} \varphi^- @ (w, \gamma_2)$

Proof (sketch): By induction over the structure of φ . When $\varphi = p$, we can show that $\gamma_1 \wedge (\wedge w) @ \top \vdash (\gamma_2 \wedge (\wedge w)) @ \top$. Then we can use the CONV rule to reach the conclusion. When $\varphi^+ = \square_{\gamma} \varphi_1^+$ the proof proceeds as follows: Given any w' , we can apply I.H. on φ_1^+ and instantiate w with $(w', \bar{\gamma})$ to obtain $\varphi_1^+ @ (w', \bar{\gamma}, \bar{\gamma}_2) \vdash \varphi_1^+ @ (w', \bar{\gamma}, \bar{\gamma}_1)$. Then we can apply $\square R$ and $\square L$ rules to reach the conclusion. \square

It is not readily apparent that Axiom Mon follows from Lemma 15. Becker et al. [5] explains that Axiom Mon says that submitting more or strengthening policies in the counterfactual statement $\square_{\gamma} \varphi$ should not make the formula φ false. This is our definition of positive formulas from Section 4, so Lemma 15 is, in fact, a more general statement which entails Axiom Mon.