

Dynamic Enforcement of Knowledge-based Security Policies[†]

Piotr Mardziel, Stephen Magill, Michael Hicks
University of Maryland, College Park

Mudhakar Srivatsa
IBM T.J. Watson Research Laboratory

Abstract—This paper explores the idea of *knowledge-based security policies*, which are used to decide whether to answer queries over secret data based on an estimation of the querier’s (possibly increased) knowledge given the results. Limiting knowledge is the goal of existing information release policies that employ mechanisms such as noising, anonymization, and redaction. Knowledge-based policies are more general: they increase flexibility by not fixing the means to restrict information flow. We enforce a knowledge-based policy by explicitly tracking a model of a querier’s belief about secret data, represented as a probability distribution, and denying any query that could increase knowledge above a given threshold. We implement query analysis and belief tracking via abstract interpretation using a novel *probabilistic polyhedral* domain, whose design permits trading off precision with performance while ensuring estimates of a querier’s knowledge are sound. Experiments with our implementation show that several useful queries can be handled efficiently, and performance scales far better than would more standard implementations of probabilistic computation based on sampling.

I. INTRODUCTION

Facebook, Twitter, Flickr, and other successful on-line services enable users to easily foster and maintain relationships by sharing information with friends and fans. These services store users’ personal information and use it to customize the user experience and to generate revenue. For example, Facebook third-party applications are granted access to a user’s “basic” data (which includes name, profile picture, gender, networks, user ID, and list of friends [1]) to implement services like birthday announcements and horoscopes, while Facebook selects ads based on age, gender, and even sexual preference [2]. Unfortunately, once personal information is collected, users have limited control over how it is used. For example, Facebook’s EULA grants Facebook a non-exclusive license to any content a user posts [3]. MySpace, another social network site, has recently begun to sell its users’ data [4].

Some researchers have proposed that, to keep tighter control over their data, users could use a storage server (e.g., running on their home network) that handles personal

data requests, and only responds when a request is deemed safe [5], [6]. The question is: which requests are safe? While deferring to user-defined access control policies seems an obvious approach, such policies are unnecessarily restrictive when the goal is to maximize the customized personal experience. To see why, consider two example applications: a horoscope or “happy birthday” application that operates on birth month and day, and a music recommendation algorithm that considers birth year (age). Access control at the granularity of the entire birth date could preclude both of these applications, while choosing only to release birth year or birth day precludes access to one application or the other. But in fact the user may not care much about these particular bits of information, but rather about what can be deduced from them. For example, it has been reported that zip code, birth date, and gender are sufficient information to uniquely identify 63% of Americans in the 2000 U.S. census [7]. So the user may be perfectly happy to reveal any one of these bits of information in its entirety as long as a querier gains no better than a $1/n$ chance to guess the entire group, for some parameter n .

This paper explores the design and implementation for enforcing what we call *knowledge-based security policies*. In our model, a user U ’s agent responds to queries involving secret data. For each querying principal Q , the agent maintains a probability distribution over U ’s secret data, representing Q ’s *belief* of the data’s likely values. For example, to mediate queries from a social networking site X , user U ’s agent may model X ’s otherwise uninformed knowledge of U ’s birthday according to a likely demographic: the birth month and day are uniformly distributed, while the birth year is most likely between 1956 and 1992 [8]. Each querier Q is also assigned a knowledge-based policy, expressed as a set of thresholds, each applying to a different group of (potentially overlapping) data. For example, U ’s policy for X might be a threshold of $1/100$ for the entire tuple (*birthdate*, *zipcode*, *gender*), and $1/5$ for just birth date. U ’s agent refuses any queries that it determines could increase Q ’s ability to guess a secret above the assigned threshold. If deemed safe, U ’s agent returns the query’s (exact) result and updates Q ’s modeled belief appropriately. (We touch upon the risk of colluding queriers shortly.)

To implement our model, we need (1) an algorithm to check whether answering a query could violate a knowledge-based policy, (2) a method for revising a querier’s belief

[†] University of Maryland, Department of Computer Science Technical Report CS-TR-4978. This paper is an extended version of the paper of the same title appearing in the proceedings of the 24th IEEE Computer Security Foundations Symposium. This version contains additional discussion (Appendices A and B, Section VI-B), and bits throughout, new and revised performance experiments (Section VII and Appendix C), and proofs of theorems (Appendices D and E).

according to the answer that is given, and (3) means to implement (1) and (2) efficiently. We build on the work of Clarkson et al. [9] (reviewed in Section III), which works out the theoretical basis for (2). The main contributions of this paper, therefore, in addition to the idea of knowledge-based policies, are our solutions to problems (1) and (3).

Given a means to revise querier beliefs based on prior answers, it seems obvious how to check that a query does not reveal too much: U runs the query, tentatively revises Q 's belief based on the result, and then responds with the answer only if Q 's revised belief about the secrets does not exceed the prescribed thresholds. Unfortunately, with this approach the decision to deny depends on the actual secret, so a rejection could leak information. We give an example in the next section that shows how the entire secret could be revealed. Therefore, we propose that a query should be rejected if there exists *any* possible secret value that could induce an output whereby the revised belief would exceed the threshold. This idea is described in detail in Section IV.

To implement belief tracking and revision, our first thought was to use languages for probabilistic computation and conditioning, which provide the foundational elements of the approach. Languages we know of—IBAL [10], Probabilistic Scheme [11], and several other systems [12], [13], [14]—are implemented using sampling. Unfortunately, we found these implementations to be inadequate because they either underestimate the querier's knowledge when sampling too little, or run too slowly when the state space is large.

Instead of using sampling, we have developed an implementation based on abstract interpretation. In Section V we develop a novel abstract domain of *probabilistic polyhedra*, which extends the standard convex polyhedron abstract domain [15] with measures of probability. We represent beliefs as a set of probabilistic (as developed in Section VI). While some prior work has explored probabilistic abstract interpretation [16], this work does not support belief revision, which is required to track how observation of outputs affects a querier's belief. Support for revision requires that we maintain both under- and over-approximations of the querier's belief, whereas [16] deals only with over-approximation. We have developed an implementation of our approach based on Parma [17] and LattE [18], which we present in Section VII along with some experimental measurements of its performance. We find that while the performance of Probabilistic Scheme degrades significantly as the input space grows, our implementation scales much better, and can be orders of magnitude faster.

Knowledge-based policies aim to ensure that an attacker's knowledge of a secret does not increase much when learning the result of a query. Much prior work aims to enforce similar properties by tracking information leakage quantitatively [19], [20], [21], [22], [23]. Our approach is more precise (but also more resource-intensive) because it maintains an on-line model of adversary knowledge. An alternative to

knowledge-based privacy is *differential privacy* [24] (DP), which requires that a query over a database of individuals' records produces roughly the same answer whether a particular individual's data is in the database or not—the possible knowledge of the querier, and the impact of the query's result on it, need not be directly considered. As such, DP avoids the danger of mismodeling a querier's knowledge and as a result inappropriately releasing information. DP also ensures a high degree of compositionality, which provides some assurance against collusion. However, DP applies once an individual has released his personal data to a trusted third party's database, a release we are motivated to avoid. Moreover, applying DP to queries over an individual's data, rather than a population, introduces so much noise that the results are often useless. We discuss these issues along with other related work in Section VIII.

The next section presents a technical overview of the rest of the paper, whose main results are contained in Sections III–VII, with further discussion and ideas for future work in Sections VIII and IX.

II. OVERVIEW

Knowledge-based policies and beliefs. User Bob would like to enforce a knowledge-based policy on his data so that advertisers do not learn too much about him. Suppose Bob considers his birthday of September 27, 1980 to be relatively private; variable $bday$ stores the calendar day (a number between 0 and 364, which for Bob would be 270) and $byear$ stores the birth year (which would be 1980). To $bday$ he assigns a *knowledge threshold* $t_d = 0.2$ stating that he does not want an advertiser to have better than a 20% likelihood of guessing his birth day. To the pair $(bday, byear)$ he assigns a threshold $t_{dy} = 0.05$, meaning he does not want an advertiser to be able to guess the combination of birth day *and* year together with better than a 5% likelihood.

Bob runs an agent program to answer queries about his data on his behalf. This agent models an estimated *belief* of queriers as a probability distribution δ , which is conceptually a map from secret states to positive real numbers representing probabilities (in range $[0, 1]$). Bob's secret state is the pair $(bday = 270, byear = 1980)$. The agent represents a distribution as a set of probabilistic polyhedra. For now, we can think of a probabilistic polyhedron as a standard convex polyhedron C with a probability mass m , where the probability of each integer point contained in C is $m/\#(C)$, where $\#(C)$ is the number of integer points contained in the polyhedron C . Shortly we present a more involved representation.

Initially, the agent might model an advertiser X 's belief using the following rectangular polyhedron C , where each point contained in it is considered equally likely ($m = 1$):

$$C = 0 \leq bday < 365, 1956 \leq byear < 1993$$

Enforcing knowledge-based policies safely. Suppose X wants to identify users whose birthday falls within the next week, to promote a special offer. X sends Bob’s agent the following program.

Example 1.

```

today := 260;
if bday ≥ today ∧ bday < (today + 7) then
  output := True;

```

This program refers to Bob’s secret variable $bday$, and also uses non-secret variables $today$, which represents the current day and is here set to be 260, and $output$, which is set to True if the user’s birthday is within the next seven days (we assume $output$ is initially False).

The agent must decide whether returning the result of running this program will potentially increase X ’s knowledge about Bob’s data above the prescribed threshold. We explain how it makes this determination shortly, but for the present we can see that answering the query is safe: the returned $output$ variable will be False which essentially teaches the querier that Bob’s birthday is not within the next week, which still leaves many possibilities. As such, the agent revises his model of the querier’s belief to be the following pair of rectangular polyhedra C_1, C_2 , where again all points in each are equally likely ($m_1 \approx 0.726, m_2 \approx 0.274$):

$$\begin{aligned}
C_1 &= 0 \leq bday < 260, 1956 \leq byear < 1993 \\
C_2 &= 267 \leq bday < 365, 1956 \leq byear < 1993
\end{aligned}$$

Ignoring $byear$, there are 358 possible values for $bday$ and each is equally likely. Thus the probability of any one is $1/358 \approx 0.0028 \leq t_d = 0.2$.

Suppose the next day the same advertiser sends the same program to Bob’s user agent, but with $today$ set to 261. Should the agent run the program? At first glance, doing so seems OK. The program will return False, and the revised belief will be the same as above but with constraint $bday \geq 267$ changed to $bday \geq 268$, meaning there is still only a $1/357 = 0.0028$ chance to guess $bday$.

But suppose Bob’s birth day was actually 267, rather than 270. The first query would have produced the same revised belief as before, but since the second query would return True (since $bday = 267 < (261+7)$), the querier can deduce Bob’s birth day exactly: $bday \geq 267$ (from the first query) and $bday < 268$ (from the second query) together imply that $bday = 267$! But the user agent is now stuck: it cannot simply refuse to answer the query, because the querier knows that with $t_d = 0.2$ (or indeed, any reasonable threshold) the only good reason to refuse is when $bday = 267$. As such, refusal essentially tells the querier the answer.

The lesson is that the decision to refuse a query must not be based on the effect of running the query on the actual secret, because then a refusal could leak information. In Section IV we propose that an agent should reject a program if there exists *any* possible secret that could cause a program




Figure 1. Example 2: most precise revised beliefs

answer to increase querier knowledge above the threshold. As such we would reject the second query regardless of whether $bday = 270$ or $bday = 267$.

Full probabilistic polyhedra. Now suppose, having run the first query and rejected the second, the user agent receives the following program from X .

Example 2.

```

age := 2011 - byear;
if age = 20 ∨ ... ∨ age = 60 then
  output := True;
pif 0.1 then output := True;

```

This program attempts to discover whether this year is a “special” year for the given user, who thus deserves a special offer. The program returns True if either the user’s age is (or will be) an exact decade, or if the user wins the luck of the draw (one chance in ten), as implemented by the probabilistic if statement.

Running this program reveals nothing about $bday$, but does reveal something about $byear$. In particular, if $output = False$ then the querier knows that $byear \notin \{1991, 1981, 1971, 1961\}$, but all other years are equally likely. We could represent this new knowledge, combined with the knowledge gained from the first query, as shown in Figure 1(a), where each shaded box is a polyhedron containing equally likely points. On the other hand, if $output = True$ then either $byear \in \{1991, 1981, 1971, 1961\}$ or the user got lucky. We represent the querier’s knowledge in this case as in Figure 1(b). Darker shading indicates higher probability; thus, all years are still possible, though some are much more likely than others. With the given threshold of $t_{dy} = 0.05$, the agent will permit the query; when $output = False$, the likelihood of any point in the shaded region is $1/11814$; when $output = True$, the points in the dark bands are the most likely, with probability $5/13067$. Since both outcomes are possible with Bob’s $byear = 1980$, the revised belief will depend on the result of the probabilistic if statement.

This example illustrates a potential problem with the simple representation of probabilistic polyhedra mentioned earlier: when $output = False$ we will jump from using two

probabilistic polyhedra to ten, and when *output* = True we jump to using eighteen. Allowing the number of polyhedra to grow without bound will result in performance problems. To address this concern, we need a way to abstract our belief representation to be more concise. Section V shows how to represent a probabilistic polyhedron P as a seven-tuple, $(C, s^{\min}, s^{\max}, p^{\min}, p^{\max}, m^{\min}, m^{\max})$ where s^{\min} and s^{\max} are lower and upper bounds on the number of points with non-zero probability in the polyhedron C (called the *support points* of C); the quantities p^{\min} and p^{\max} are lower and upper bounds on the probability mass per support point; and m^{\min} and m^{\max} give bounds on the total probability mass. Thus, polyhedra modeled using the simpler representation (C, m) given earlier are equivalent to ones in the more involved representation with $m^{\max} = m^{\min} = m$, $p^{\max} = p^{\min} = m/\#(C)$, and $s^{\max} = s^{\min} = \#(C)$.

With this representation, we could choose to collapse the sets of polyhedron given in Figure 1. For example, we could represent Figure 1(a) with two probabilistic polyhedra P_1 and P_2 containing polyhedra C_1 and C_2 defined above, respectively, essentially drawing a box around the two groupings of smaller boxes in the figure. The other parameters for P_1 would be as follows:

$$\begin{aligned} p_1^{\min} &= p_1^{\max} = 9/135050 \\ s_1^{\min} &= s_1^{\max} = 8580 \\ m_1^{\min} &= m_1^{\max} = 7722/13505 \end{aligned}$$

Notice that $s_1^{\min} = s_2^{\max} = 8580 < \#(C_1) = 9620$, illustrating that the ‘‘bounding box’’ of the polyhedron covers more area than is strictly necessary. In this representation the probabilities may not be normalized, which improves both performance and precision. For this example, P_2 happens to have $m_2^{\min} = m_2^{\max} = 14553/67525$ so we can see $m_1^{\max} + m_2^{\max} = (53163/67525) \neq 1$.

If we consider the representation of Figure 1(b) in a similar manner, using the same two polyhedra C_1 and C_2 , the other parameters for C_1 are as follows:

$$\begin{aligned} p_1^{\min} &= 1/135050 & p_1^{\max} &= 10/135050 \\ s_1^{\min} &= 9620 & s_1^{\max} &= 9620 \\ m_1^{\min} &= 26/185 & m_1^{\max} &= 26/185 \end{aligned}$$

In this case $s_1^{\min} = s_1^{\max} = \#(C_1)$, meaning that all covered points are possible, but $p_1^{\min} \neq p_1^{\max}$ as some points are more probable than others (i.e., those in the darker band).

The key property of probabilistic polyhedra, and a main technical contribution of this paper, is that this abstraction can be used to make sound security policy decisions. To accept a query, we must check that, for all possible outputs, the querier’s revised, normalized belief of any of the possible secrets is below the threshold t . In checking whether the revised beliefs in our example are acceptable, the agent will try to find the maximum probability the querier could ascribe to a state, for each possible output. In the case *output* = True, the most probable points are those in the dark bands,

<i>Variables</i>	x	\in	Var
<i>Integers</i>	n	\in	\mathbb{Z}
<i>Rationals</i>	q	\in	\mathbb{Q}
<i>Arith.ops</i>	aop	$::=$	$+ \mid \times \mid -$
<i>Rel.ops</i>	$relop$	$::=$	$\leq \mid < \mid = \mid \neq \mid \dots$
<i>Arith.exps</i>	E	$::=$	$x \mid n \mid E_1 \ aop \ E_2$
<i>Bool.exps</i>	B	$::=$	$E_1 \ relop \ E_2 \mid$ $B_1 \wedge B_2 \mid B_1 \vee B_2 \mid \neg B$
<i>Statements</i>	S	$::=$	skip $\mid x := E \mid$ if B then S_1 else $S_2 \mid$ pif q then S_1 else $S_2 \mid$ $S_1 ; S_2 \mid$ while B do S

Figure 2. Core language syntax

which each have probability mass $10/135050 = p_1^{\max}$ (the dark bands in P_2 have the same probability). To find the maximum normalized probability of these points, we divide by the minimum possible total mass, as given by the lower bounds in our abstraction. In our example, this results in $p_1^{\max}/(m_1^{\min} + m_2^{\min}) = (10/135050)/(26/185 + 49/925) \approx 0.0004 \leq t_d = 0.05$.

As just shown, the bound on minimum total mass is needed in order to soundly normalize distributions in our abstraction. The maintenance of such lower bounds on probability mass is a key component of our abstraction that is missing from prior work. Each of the components of a probabilistic polyhedron play a role in producing the lower bound on total mass. While $s_1^{\min}, s_1^{\max}, p_1^{\min}$, and m_1^{\max} do not play a role in making the final policy decision, their existence allows us to more accurately update belief during the query evaluation that precedes the final policy check. The choice of the number of probabilistic polyhedra to use impacts both precision and performance, so choosing the right number is a challenge. For the examples given in this section, our implementation can often answer queries in a few seconds; details are in Sections V–VII.

III. TRACKING BELIEFS

This section reviews Clarkson et al.’s method of revising a querier’s belief of the possible valuations of secret variables based on the result of a query involving those variables [9].

A. Core language

The programming language we use for queries is given in Figure 2. A computation is defined by a statement S whose standard semantics can be viewed as a relation between states: given an input state σ , running the program will produce an output state σ' . States are maps from variables to integers:

$$\sigma, \tau \in \mathbf{State} \stackrel{\text{def}}{=} \mathbf{Var} \rightarrow \mathbb{Z}$$

Sometimes we consider states with domains restricted to a subset of variables V , in which case we write $\sigma_V \in$

$\mathbf{State}_V \stackrel{\text{def}}{=} V \rightarrow \mathbb{Z}$. We may also *project* states to a set of variables V :

$$\sigma \upharpoonright V \stackrel{\text{def}}{=} \lambda x \in \mathbf{Var}_V. \sigma(x)$$

The language is essentially standard. We limit the form of expressions to support our abstract interpretation-based semantics (Section V). The semantics of the statement form `pif q then S_1 else S_2` is non-deterministic: the result is that of S_1 with probability q , and S_2 with probability $1 - q$.

B. Probabilistic semantics for tracking beliefs

To enforce a knowledge-based policy, a user agent must be able to estimate what a querier could learn from the output of his query. To do this, the agent keeps a distribution δ that represents the querier’s *belief* of the likely valuations of the user’s secrets. More precisely, a distribution is a map from states to positive real numbers, interpreted as probabilities (in range $[0, 1]$).

$$\delta \in \mathbf{Dist} \stackrel{\text{def}}{=} \mathbf{State} \rightarrow \mathbb{R}_+$$

We sometimes focus our attention on distributions over states of a fixed set of variables V , in which case we write $\delta_V \in \mathbf{Dist}_V$ to mean $\mathbf{State}_V \rightarrow \mathbb{R}_+$. Projecting distributions onto a set of variables is as follows:¹

$$\delta \upharpoonright V \stackrel{\text{def}}{=} \lambda \sigma_V \in \mathbf{State}_V. \sum_{\sigma' \upharpoonright V = \sigma_V} \delta(\sigma')$$

The *mass* of a distribution, written $\|\delta\|$ is the sum of the probabilities ascribed to states, $\sum_{\sigma} \delta(\sigma)$. A *normalized distribution* is one such that $\|\delta\| = 1$. A normalized distribution can be constructed by scaling a distribution according to its mass:

$$\text{normal}(\delta) \stackrel{\text{def}}{=} \frac{1}{\|\delta\|} \cdot \delta$$

The *support* of a distribution is the set of states which have non-zero probability: $\text{support}(\delta) \stackrel{\text{def}}{=} \{\sigma \mid \delta(\sigma) > 0\}$.

The agent evaluates a query in light of the querier’s initial belief using a probabilistic semantics. Figure 3 defines a semantic function $\llbracket \cdot \rrbracket$ whereby $\llbracket S \rrbracket \delta = \delta'$ indicates that, given an input distribution δ , the semantics of program S is the output distribution δ' . The semantics is defined in terms of operations on distributions, including *assignment* $\delta[v \rightarrow E]$ (used in the rule for `$v := E$`), *conditioning* $\delta|B$ and *addition* $\delta_1 + \delta_2$ (used in the rule for `if`), and *scaling* $q \cdot \delta$ where q is a rational (used for `pif`). The semantics is standard (cf. Clarkson et al. [9]). A brief review is given in Appendix A.

¹The notation $\sum_{x \mid \pi} \rho$ can be read ρ is the sum over all x such that formula π is satisfied (where x is bound in ρ and π).

$$\begin{aligned} \llbracket \text{skip} \rrbracket \delta &= \delta \\ \llbracket x := E \rrbracket \delta &= \delta[x \rightarrow E] \\ \llbracket \text{if } B \text{ then } S_1 \text{ else } S_2 \rrbracket \delta &= \llbracket S_1 \rrbracket (\delta|B) + \llbracket S_2 \rrbracket (\delta|\neg B) \\ \llbracket \text{pif } q \text{ then } S_1 \text{ else } S_2 \rrbracket \delta &= \llbracket S_1 \rrbracket (q \cdot \delta) + \llbracket S_2 \rrbracket ((1 - q) \cdot \delta) \\ \llbracket S_1 ; S_2 \rrbracket \delta &= \llbracket S_2 \rrbracket (\llbracket S_1 \rrbracket \delta) \\ \llbracket \text{while } B \text{ do } S \rrbracket &= \text{lfp} [\lambda f : \mathbf{Dist} \rightarrow \mathbf{Dist}. \lambda \delta. \\ &\quad f(\llbracket S \rrbracket (\delta|B)) + (\delta|\neg B)] \end{aligned}$$

where

$$\begin{aligned} \delta[x \rightarrow E] &\stackrel{\text{def}}{=} \lambda \sigma. \sum_{\tau \mid \tau[x \rightarrow [E]\tau] = \sigma} \delta(\tau) \\ \delta_1 + \delta_2 &\stackrel{\text{def}}{=} \lambda \sigma. \delta_1(\sigma) + \delta_2(\sigma) \\ \delta|B &\stackrel{\text{def}}{=} \lambda \sigma. \text{if } \llbracket B \rrbracket \sigma \text{ then } \delta(\sigma) \text{ else } 0 \\ p \cdot \delta &\stackrel{\text{def}}{=} \lambda \sigma. p \cdot \delta(\sigma) \end{aligned}$$

Figure 3. Probabilistic semantics for the core language

C. Belief and security

Clarkson et al. [9] describe how a belief about possible values of a secret, expressed as a probability distribution, can be revised according to an experiment using the actual secret. Such an experiment works as follows.

The values of the set of secret variables H are given by the hidden state σ_H . The attacker’s initial belief as to the possible values of σ_H is represented as a distribution δ_H . A query is a program S that makes use of variables H and possibly other, non-secret variables from a set L ; the final values of L , after running S , are made visible to the attacker. Let σ_L be an arbitrary initial state of these variables such that $\text{domain}(\sigma_L) = L$. Then we take the following steps:

Step 1. Evaluate S probabilistically using the attacker’s belief about the secret to produce an output distribution δ' , which amounts to the attacker’s prediction of the possible output states. This is computed as $\delta' = \llbracket S \rrbracket \delta$, where δ , a distribution over variables $H \uplus L$, is defined as $\delta = \delta_H \times \dot{\sigma}_L$. Here, we make use of the distribution product operator and point operator. That is, given δ_1, δ_2 , which are distributions over states having disjoint domains, the *distribution product* is

$$\delta_1 \times \delta_2 \stackrel{\text{def}}{=} \lambda(\sigma_1, \sigma_2). \delta_1(\sigma_1) \cdot \delta_2(\sigma_2)$$

where (σ_1, σ_2) is the “concatenation” of the two states, which is itself a state and is well-defined because the two states’ domains are disjoint. And, given a state σ , the *point distribution* $\dot{\sigma}$ is a distribution in which only σ is possible:

$$\dot{\sigma} \stackrel{\text{def}}{=} \lambda \tau. \text{if } \sigma = \tau \text{ then } 1 \text{ else } 0$$

Thus, the initial distribution δ is the attacker’s belief about the secret variables combined with an arbitrary valuation of the public variables.

Step 2. Using the actual secret σ_H , evaluate S “concretely” to produce an output state $\hat{\sigma}_L$, in three steps. First, we have $\hat{\delta}' = \llbracket S \rrbracket \hat{\delta}$, where $\hat{\delta} = \dot{\sigma}_H \times \dot{\sigma}_L$. Second, we have $\hat{\sigma} \in \Gamma(\hat{\delta})$ where Γ is a sampling operator that produces a

state σ from the domain of a distribution δ with probability $\delta(\sigma)/\|\delta\|$. Finally, we extract the attacker-visible output of the sampled state by projecting away the high variables: $\hat{\sigma}_L = \hat{\sigma} \upharpoonright L$.

Step 3. Revise the attacker’s initial belief δ_H according to the observed output $\hat{\sigma}_L$, yielding a new belief $\hat{\delta}_H = \delta' | \hat{\sigma}_L \upharpoonright H$. Here, δ' is *conditioned* on the output $\hat{\sigma}_L$, which yields a new distribution, and this distribution is then projected to the variables H . The conditioning operation is defined as follows:

$$\delta | \sigma_V \stackrel{\text{def}}{=} \lambda \sigma. \text{ if } \sigma \upharpoonright V = \sigma_V \text{ then } \delta(\sigma) \text{ else } 0$$

Note that this protocol assumes that S always terminates and does not modify the secret state. The latter assumption can be eliminated by essentially making a copy of the state before running the program, while eliminating the former depends on the observer’s ability to detect nontermination [9].

IV. ENFORCING KNOWLEDGE-BASED POLICIES

When presented with a query over a user’s data σ_H , the user’s agent should only answer the query if doing so will not reveal too much information. More precisely, given a query S , the agent will return the public output σ_L resulting from running S on σ_H if the agent deems that from this output the querier cannot guess the secret state σ_H beyond some level of doubt, identified by a threshold t . If this threshold could be exceeded, then the agent declines to run S . We call this security check *knowledge threshold security*.

Definition 3 (Knowledge Threshold Security). Let $\delta' = \llbracket S \rrbracket \delta$, where δ is the model of the querier’s initial belief. Then query S is *threshold secure* iff for all $\sigma_L \in \text{support}(\delta' \upharpoonright L)$ and all $\sigma'_H \in \mathbf{State}_H$ we have $(\text{normal}((\delta' | \sigma_L) \upharpoonright H))(\sigma'_H) \leq t$ for some threshold t .

This definition can be related to the experiment protocol defined in Section III-C. First, δ' in the definition is the same as δ' computed in the first step of the protocol. Step 2 in the protocol produces a concrete output $\hat{\sigma}_L$ based on executing S on the actual secret σ_H , and Step 3 revises the querier’s belief based on this output. Definition 3 generalizes these two steps: instead of considering a single concrete output based on the actual secret it considers *all possible* concrete outputs, as given by $\text{support}(\delta' \upharpoonright L)$, and ensures that the revised belief in each case for *all possible* secret states must assign probability no greater than t .

This definition considers a threshold for the whole secret state σ_H . As described in Section II we can also enforce thresholds over portions of a secret state. In particular, a threshold that applies only to variables $V \subseteq H$ requires that all $\sigma'_V \in \mathbf{State}_V$ result in $(\text{normal}(\delta' | \sigma_L \upharpoonright V))(\sigma'_V) \leq t$.

The two “foralls” in the definition are critical for ensuring security. The reason was shown by the first example in Section II: If we used the flawed approach of just running the experiment protocol and checking if $\hat{\delta}_H(\sigma_H) > t$

then rejection depends on the value of the secret state and could reveal information about it. The more general policy $\forall \sigma_L \in \text{support}(\delta' \upharpoonright L). (\text{normal}(\delta' | \sigma_L \upharpoonright H))(\sigma_H) \leq t$, would sidestep the problem in the example, but this policy could still reveal information because it, too, depends on the actual secret σ_H . (An example illustrating the problem in this case is given in Appendix B.) Definition 3 avoids any inadvertent information leakage because rejection is not based on the actual secret: if there exists *any* secret such that a possible output would reveal too much, the query is rejected. Definition 3 resembles, but is stronger than, *min-entropy*, as the security decision is based on the most likely secret from the attacker’s point of view [20]; further details are given in Section VIII.

V. BELIEF REVISION VIA ABSTRACT INTERPRETATION

Consider how we might implement belief tracking and revision to enforce the threshold security property given in Definition 3. A natural choice would be to evaluate queries using a probabilistic programming language with support for conditioning; examples are IBAL [10], Probabilistic Scheme [11], and several others [12], [13], [14]. In these languages, probabilistic evaluation is achieved by enumerating inputs (sampling). Probabilities are associated with each input and tracked during execution. As more inputs are enumerated, a more complete view of the output distribution emerges. Unfortunately, to get an accurate estimate of the revised distribution following an output observation, one must enumerate the entire input space, which could be quite large. If insufficient coverage is achieved, then the threshold check in Definition 3 could either be unsound or excessively conservative, depending in which direction an implementation errs.

To avoid sampling, we have developed a new means to perform probabilistic computation based on abstract interpretation. In this approach, execution time depends on the complexity of the query rather than the size of the input space. In the next two sections, we present two abstract domains. This section presents the first, denoted \mathbb{P} , where an abstract element is a single *probabilistic polyhedron*, which is a convex polyhedron [15] with information about the probabilities of its points. Because using a single polyhedron will accumulate imprecision after multiple queries, in our implementation we actually use a different domain, denoted $\mathcal{P}_n(\mathbb{P})$, for which an abstract element consists of a set of at most n probabilistic polyhedra (whose construction is inspired by powersets of polyhedra [?], [?]). This domain, described in the next section, allows us to retain precision at the cost of increased execution time. By adjusting n , the user can trade off efficiency and precision.

A. Polyhedra

We first review *convex polyhedra*, a common technique for representing sets of program states. We use the meta-

variables β, β_1, β_2 , etc. to denote linear inequalities. We write $fv(\beta)$ to be the set of variables occurring in β ; we also extend this to sets, writing $fv(\{\beta_1, \dots, \beta_n\})$ for $fv(\beta_1) \cup \dots \cup fv(\beta_n)$.

Definition 4. A *convex polyhedron* $C = (B, V)$ is a set of linear inequalities $B = \{\beta_1, \dots, \beta_m\}$, interpreted conjunctively, over dimensions V . We write \mathbb{C} for the set of all convex polyhedra. A polyhedron C represents a set of states, denoted $\gamma_{\mathbb{C}}(C)$, as follows, where $\sigma \models \beta$ indicates that the state σ satisfies the inequality β .

$$\gamma_{\mathbb{C}}((B, V)) \stackrel{\text{def}}{=} \{\sigma \mid \text{domain}(\sigma) = V, \forall \beta \in B. \sigma \models \beta\}$$

Naturally we require that $fv(\{\beta_1, \dots, \beta_n\}) \subseteq V$. We write $fv((B, V))$ to denote the set of variables V of a polyhedron.

Given a state σ and an ordering on the variables in $\text{domain}(\sigma)$, we can view σ as a point in an N -dimensional space, where $N = |\text{domain}(\sigma)|$. The set $\gamma_{\mathbb{C}}(C)$ can then be viewed as the integer-valued lattice points in an N -dimensional polyhedron. Due to this correspondence, we use the words *point* and *state* interchangeably. We will sometimes write linear equalities $x = f(\vec{y})$ as an abbreviation for the pair of inequalities $x \leq f(\vec{y})$ and $x \geq f(\vec{y})$.

Let $C = (B, V)$. Convex polyhedra support the following operations.

- Polyhedron size, or $\#(C)$, is the number of integer points in the polyhedron, i.e., $|\gamma_{\mathbb{C}}(C)|$. We will always consider bounded polyhedra when determining their size, ensuring that $\#(C)$ is finite.
- Expression evaluation, $\langle\langle B \rangle\rangle C$ returns a convex polyhedron containing at least the points in C that satisfy B .
- Expression count, $C \# B$ returns an upper bound on the number of integer points in C that satisfy B . (It may be more precise than $\#(\langle\langle B \rangle\rangle C)$.)
- Meet, $C_1 \sqcap_{\mathbb{C}} C_2$ is the convex polyhedron containing exactly the set of points in the intersection of $\gamma_{\mathbb{C}}(C_1), \gamma_{\mathbb{C}}(C_2)$.
- Join, $C_1 \sqcup_{\mathbb{C}} C_2$ is the smallest convex polyhedron containing both $\gamma_{\mathbb{C}}(C_1)$ and $\gamma_{\mathbb{C}}(C_2)$.
- Comparison, $C_1 \sqsubseteq_{\mathbb{C}} C_2$ is a partial order whereby $C_1 \sqsubseteq_{\mathbb{C}} C_2$ if and only if $\gamma_{\mathbb{C}}(C_1) \subseteq \gamma_{\mathbb{C}}(C_2)$.
- Affine transform, $C[x \rightarrow E]$, where $x \in fv(C)$, computes an affine transformation of C . This scales the dimension corresponding to x by the coefficient of x in E and shifts the polyhedron. For example, $(\{x \leq y, y = 2z\}, V)[y \rightarrow z + y]$ evaluates to $(\{x \leq y - z, y - z = 2z\}, V)$.
- Forget, $f_x(C)$, projects away x . That is, $f_x(C) = \pi_{fv(C) - \{x\}}(C)$, where $\pi_V(C)$ is a polyhedron C' such that $\gamma_{\mathbb{C}}(C') = \{\sigma \mid \sigma' \in \gamma_{\mathbb{C}}(C) \wedge \sigma = \sigma' \upharpoonright V\}$. So $C' = f_x(C)$ implies $x \notin fv(C')$.

We write $isempty(C)$ iff $\gamma_{\mathbb{C}}(C) = \emptyset$.

B. Probabilistic Polyhedra

We take this standard representation of sets of program states and extend it to a representation for sets of distribu-

tions over program states. We define *probabilistic polyhedra*, the core element of our abstract domain, as follows.

Definition 5. A *probabilistic polyhedron* P is a tuple $(C, s^{\min}, s^{\max}, p^{\min}, p^{\max}, m^{\min}, m^{\max})$. We write \mathbb{P} for the set of probabilistic polyhedra. The quantities s^{\min} and s^{\max} are lower and upper bounds on the number of support points in the polyhedron C . The quantities p^{\min} and p^{\max} are lower and upper bounds on the probability mass *per support point*. The m^{\min} and m^{\max} components give bounds on the total probability mass. Thus P represents the *set* of distributions $\gamma_{\mathbb{P}}(P)$ defined below.

$$\begin{aligned} \gamma_{\mathbb{P}}(P) \stackrel{\text{def}}{=} \{ & \delta \mid \text{support}(\delta) \subseteq \gamma_{\mathbb{C}}(C) \wedge \\ & s^{\min} \leq |\text{support}(\delta)| \leq s^{\max} \wedge \\ & m^{\min} \leq \|\delta\| \leq m^{\max} \wedge \\ & \forall \sigma \in \text{support}(\delta). p^{\min} \leq \delta(\sigma) \leq p^{\max} \} \end{aligned}$$

We will write $fv(P) \stackrel{\text{def}}{=} fv(C)$ to denote the set of variables used in the probabilistic polyhedron.

Note the set $\gamma_{\mathbb{P}}(P)$ is singleton exactly when $s^{\min} = s^{\max} = \#(C)$ and $p^{\min} = p^{\max}$, and $m^{\min} = m^{\max}$. In such a case $\gamma_{\mathbb{P}}(P)$ is the uniform distribution where each state in $\gamma_{\mathbb{C}}(C)$ has probability p^{\min} . Distributions represented by a probabilistic polyhedron are not necessarily normalized (as was true in Section III-B). In general, there is a relationship between p^{\min}, s^{\min} , and m^{\min} , in that $m^{\min} \geq p^{\min} \cdot s^{\min}$ (and $m^{\max} \leq p^{\max} \cdot s^{\max}$), and the combination of the three can yield more information than any two in isolation.

Our convention will be to use $C_1, s_1^{\min}, s_1^{\max}$, etc. for the components associated with probabilistic polyhedron P_1 and to use subscripts to name different probabilistic polyhedra.

Distributions are ordered point-wise [9]. That is, $\delta_1 \leq \delta_2$ if and only if $\forall \sigma. \delta_1(\sigma) \leq \delta_2(\sigma)$. For our abstract domain, we say that $P_1 \sqsubseteq_{\mathbb{P}} P_2$ if and only if $\forall \delta_1 \in \gamma_{\mathbb{P}}(P_1). \exists \delta_2 \in \gamma_{\mathbb{P}}(P_2). \delta_1 \leq \delta_2$. Testing $P_1 \sqsubseteq_{\mathbb{P}} P_2$ mechanically is non-trivial, but is unnecessary in our semantics. Rather, we need to test whether a distribution represents only the zero distribution $0_{\text{Dist}} \stackrel{\text{def}}{=} \lambda \sigma. 0$ in order to see that a fixed point for evaluating $\langle\langle \text{while } B \text{ do } S \rangle\rangle P$ has been reached. Intuitively, no further iterations of the loop need to be considered once the probability mass flowing into the n^{th} iteration is zero. This condition can be detected as follows:

$$\begin{aligned} iszero(P) \stackrel{\text{def}}{=} & \\ & s^{\min} = s^{\max} = 0 \wedge m^{\min} = 0 \leq m^{\max} \\ & \vee m^{\min} = m^{\max} = 0 \wedge s^{\min} = 0 \leq s^{\max} \\ & \vee isempty(C) \wedge s^{\min} = 0 \leq s^{\max} \wedge m^{\min} = 0 \leq m^{\max} \\ & \vee p^{\min} = p^{\max} = 0 \wedge s^{\min} = 0 \leq s^{\max} \wedge m^{\min} = 0 \leq m^{\max} \end{aligned}$$

If $iszero(P)$ holds, it is the case that $\gamma_{\mathbb{P}}(P) = \{0_{\text{Dist}}\}$. Note that having a more conservative definition of this function (which holds for fewer probabilistic polyhedra) would be reasonable since it would simply mean our analysis would terminate less often than it could, with no effect on security. More details are given in Appendix D.

In a standard abstract domain, termination of the fixed point computation for loops is often ensured by use of a widening operator. This allows abstract fixed points to be computed in fewer iterations and also permits analysis of loops that may not terminate. In our setting, non-termination may reveal information about secret values. As such, we would like to reject queries that may be non-terminating.

We enforce this by not introducing a widening operator. Our abstract interpretation then has the property that it will not terminate if a loop in the query may be non-terminating (and, since it is an over-approximate analysis, it may also fail to terminate even for some terminating computations). We then reject all queries for which our analysis fails to terminate. Loops do not play a major role in any of our examples, and so this approach has proved sufficient so far. We leave for future work the development of a widening operator that soundly accounts for non-termination behavior.

Following standard abstract interpretation terminology, we will refer to $\mathcal{P}(\mathbf{Dist})$ (sets of distributions) as the *concrete domain*, \mathbb{P} as the *abstract domain*, and $\gamma_{\mathbb{P}} : \mathbb{P} \rightarrow \mathcal{P}(\mathbf{Dist})$ as the *concretization function* for \mathbb{P} .

C. Abstract Semantics for \mathbb{P}

To support execution in the abstract domain just defined, we need to provide abstract implementations of the basic operations of assignment, conditioning, addition, and scaling used in the concrete semantics given in Figure 3. We will overload notation and use the same syntax for the abstract operators as we did for the concrete operators.

As we present each operation, we will also state the associated soundness theorem which shows that the abstract operation is an over-approximation of the concrete operation. Proofs are given in Appendix D. The abstract program semantics is then exactly the semantics from Figure 3, but making use of the abstract operations defined here, rather than the operations on distributions defined in Section III-B. We will write $\langle\langle S \rangle\rangle P$ to denote the result of executing S using the abstract semantics. The main soundness theorem we obtain is the following.

Theorem 6. *For all P, δ , if $\delta \in \gamma_{\mathbb{P}}(P)$ and $\langle\langle S \rangle\rangle P$ terminates, then $\llbracket S \rrbracket \delta$ terminates and $\llbracket S \rrbracket \delta \in \gamma_{\mathbb{P}}(\langle\langle S \rangle\rangle P)$.*

When we say $\llbracket S \rrbracket \delta$ terminates (or $\langle\langle S \rangle\rangle P$ terminates) we mean that only a finite number of loop unrollings are required to interpret the statement on a particular distribution (or probabilistic polyhedron). The precise definitions of termination can be found in Appendix D.

We now present the abstract operations.

1) *Forget:* We first describe the abstract forget operator $f_y(P_1)$, which is used in implementing assignment. When we forget variable y , we collapse any states that are equivalent up to the value of y into a single state. To do this correctly, we must find an upper bound h_y^{\max} and a lower bound h_y^{\min} on the number of points that share the same value of other




Figure 4. Example of a forget operation in the abstract domain \mathbb{P} . In this case, $h_y^{\min} = 1$ and $h_y^{\max} = 3$. Note that h_y^{\max} is precise while h_y^{\min} is an under-approximation. If $s_1^{\min} = s_1^{\max} = 9$ then we have $s_2^{\min} = 3$, $s_2^{\max} = 4$, $p_2^{\min} = p_1^{\min} \cdot 1$, $p_2^{\max} = p_2^{\max} \cdot 4$.

dimensions x (this may be visualized of as the min and max height of C_1 in the y dimension). Once these are obtained, we have that $f_y(P_1) \stackrel{\text{def}}{=} P_2$ where the following hold of P_2 .

$$\begin{aligned} C_2 &= f_y(C_1) \\ p_2^{\min} &= p_1^{\min} \cdot \max \{ h_y^{\min} - (\#(C_1) - s_1^{\min}), 1 \} \\ p_2^{\max} &= p_1^{\max} \cdot \min \{ h_y^{\max}, s_1^{\max} \} \\ s_2^{\min} &= \lceil s_1^{\min} / h_y^{\max} \rceil & m_2^{\min} &= m_1^{\min} \\ s_2^{\max} &= \min \{ \#(f_y(C_1)), s_1^{\max} \} & m_2^{\max} &= m_1^{\max} \end{aligned}$$

Figure 4 gives an example of a forget operation and illustrates the quantities h_y^{\max} and h_y^{\min} . If $C_1 = (B_1, V_1)$, the upper bound h_y^{\max} can be found by maximizing $y - y'$ subject to the constraints $B_1 \cup B_1[y'/y]$, where y' is a fresh variable and $B_1[y'/y]$ represents the set of constraints obtained by substituting y' for y in B_1 . As our points are integer-valued, this is an integer linear programming problem (and can be solved by ILP solvers). A less precise upper bound can be found by simply taking the extent of the polyhedron C_1 along y , which is given by $\#(\pi_y(C_1))$.

For the lower bound, it is always sound to use $h_y^{\min} = 1$, which is what our implementation does. A more precise estimate can be obtained by finding the vertex with minimal height along dimension y . Call this distance u . Since the shape is convex, all other points will have y height greater than or equal to u . We then find the smallest number of integer points that can be covered by a line segment of length u . This is given by $\lceil u \rceil - 1$. This value can be taken as h_y^{\min} .

Since the forget operator is related to projection, we state soundness in terms of the projection operation on distributions. Note that $f_v(\delta) \stackrel{\text{def}}{=} \text{domain}(\text{domain}(\delta))$, i.e., the domain of states to which δ assigns probability mass.

Lemma 7. *If $\delta \in \gamma_{\mathbb{P}}(P)$ then $\delta \upharpoonright (f_v(\delta) - \{y\}) \in \gamma_{\mathbb{P}}(f_y(P))$.*

We can define an abstract version of projection using forget:

Definition 8. Let $f_{\{x_1, x_2, \dots, x_n\}}(P) = f_{\{x_2, \dots, x_n\}}(f_{x_1}(P))$. Then $P \upharpoonright V' = f_{(\text{domain}(P) - V')}(P)$.

That is, in order to project onto the set of variables V' , we forget all variables not in V' .

2) *Assignment*: We have two cases for abstract assignment. If $x := E$ is invertible,² the result of the assignment $P_1[x \rightarrow E]$ is the probabilistic polyhedron P_2 such that $C_2 = C_1[x \rightarrow E]$ and all other components are unchanged.

If the assignment is not invertible, then information about the previous value of x is lost. In this case, we use the forget operation to project onto the other variables and then add a new constraint on x . Let $P_2 = f_x(P_1)$ where $C_2 = (B_2, V_2)$. Then $P_1[x \rightarrow E]$ is the probabilistic polyhedron P_3 with $C_3 = (B_2 \cup \{x = E\}, V_2 \cup \{x\})$ and all other components as in P_2 .

Lemma 9. *If $\delta \in \gamma_{\mathbb{P}}(P)$ then $\delta[v \rightarrow E] \in \gamma_{\mathbb{P}}(P[v \rightarrow E])$.*

The soundness of assignment relies on the fact that our language of expressions does not include division. An invariant of our representation is that $s^{\max} \leq \#(C)$. When E contains only multiplication and addition the above rules preserve this invariant; an E containing division would violate it. Division would collapse multiple points to one and so could be handled similarly to projection.

3) *Plus*: To soundly compute the effect of plus we need to determine the minimum and maximum number of points in the intersection that may be a support point for both P_1 and for P_2 . We refer to these counts as the *pessimistic overlap* and *optimistic overlap*, respectively, and define them below.

Definition 10. Given two distributions δ_1, δ_2 , we refer to the set of states that are in the support of both δ_1 and δ_2 as the *overlap* of δ_1, δ_2 . The *pessimistic overlap* of P_1 and P_2 , denoted $P_1 \odot P_2$, is the cardinality of the smallest possible overlap for any distributions $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$. The *optimistic overlap* $P_1 \ominus P_2$ is the cardinality of the largest possible overlap. Formally, we define these as follows. $n_3 \stackrel{\text{def}}{=} \#(C_1 \sqcap_C C_2)$, $n_1 \stackrel{\text{def}}{=} \#(C_1) - n_3$, and $n_2 \stackrel{\text{def}}{=} \#(C_2) - n_3$. Then

$$\begin{aligned} P_1 \odot P_2 &\stackrel{\text{def}}{=} \max \{ (s_1^{\min} - n_1) + (s_2^{\min} - n_2) - n_3, 0 \} \\ P_1 \ominus P_2 &\stackrel{\text{def}}{=} \min \{ s_1^{\max}, s_2^{\max}, n_3 \} \end{aligned}$$

We can now define abstract addition.

Definition 11. If not $iszero(P_1)$ and not $iszero(P_2)$ then $P_1 + P_2$ is the probabilistic polyhedron $P_3 =$

$(C_3, s_3^{\min}, s_3^{\max}, p_3^{\min}, p_3^{\max})$ defined as follows.

$$\begin{aligned} C_3 &= C_1 \sqcup_C C_2 \\ p_3^{\min} &= \begin{cases} p_1^{\min} + p_2^{\min} & \text{if } P_1 \odot P_2 = \#(C_3) \\ \min \{ p_1^{\min}, p_2^{\min} \} & \text{otherwise} \end{cases} \\ p_3^{\max} &= \begin{cases} p_1^{\max} + p_2^{\max} & \text{if } P_1 \odot P_2 > 0 \\ \max \{ p_1^{\max}, p_2^{\max} \} & \text{otherwise} \end{cases} \\ s_3^{\min} &= \max \{ s_1^{\min} + s_2^{\min} - P_1 \odot P_2, 0 \} \\ s_3^{\max} &= \min \{ s_1^{\max} + s_2^{\max} - P_1 \odot P_2, \#(C_3) \} \\ m_3^{\min} &= m_1^{\min} + m_2^{\min} \quad | \quad m_3^{\max} = m_1^{\max} + m_2^{\max} \end{aligned}$$

If $iszero(P_1)$ then we define $P_1 + P_2$ as identical to P_2 ; if $iszero(P_2)$, the sum is defined as identical to P_1 .

Lemma 12. *If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$ then $\delta_1 + \delta_2 \in \gamma_{\mathbb{P}}(P_1 + P_2)$.*

4) *Product*: When evaluating the product $P_3 = P_1 \times P_2$, we assume that the domains of P_1 and P_2 are disjoint, i.e., C_1 and C_2 refer to disjoint sets of variables. If $C_1 = (B_1, V_1)$ and $C_2 = (B_2, V_2)$, then the polyhedron $C_1 \times C_2 \stackrel{\text{def}}{=} (B_1 \cup B_2, V_1 \cup V_2)$ is the Cartesian product of C_1 and C_2 and contains all those states σ for which $\sigma \upharpoonright V_1 \in \gamma_C(C_1)$ and $\sigma \upharpoonright V_2 \in \gamma_C(C_2)$. Determining the remaining components is straightforward since P_1 and P_2 are disjoint.

$$\begin{aligned} C_3 &= C_1 \times C_2 \\ \begin{aligned} p_3^{\min} &= p_1^{\min} \cdot p_2^{\min} \\ s_3^{\min} &= s_1^{\min} \cdot s_2^{\min} \\ m_3^{\min} &= m_1^{\min} \cdot m_2^{\min} \end{aligned} & \left| \begin{aligned} p_3^{\max} &= p_1^{\max} \cdot p_2^{\max} \\ s_3^{\max} &= s_1^{\max} \cdot s_2^{\max} \\ m_3^{\max} &= m_1^{\max} \cdot m_2^{\max} \end{aligned} \end{aligned}$$

Lemma 13. *For all P_1, P_2 such that $fv(P_1) \cap fv(P_2) = \emptyset$, if $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$ then $\delta_1 \times \delta_2 \in \gamma_{\mathbb{P}}(P_1 \times P_2)$.*

In our examples we often find it useful to express uniformly distributed data directly, rather than encoding it using pif. In particular, consider extending statements S to include the statement form uniform $x \ n_1 \ n_2$ whose semantics is to define variable x as having values uniformly distributed between n_1 and n_2 . Its semantics is as follows.

$$\llbracket \text{uniform } x \ n_1 \ n_2 \rrbracket P_1 = f_x(P_1) \times P_2$$

Here, P_2 has $p_2^{\min} = p_2^{\max} = \frac{1}{n_2 - n_1 + 1}$, $s_2^{\min} = s_2^{\max} = n_2 - n_1 + 1$, $m_2^{\min} = m_2^{\max} = 1$, and $C_2 = (\{x \geq n_1, x \leq n_2\}, \{x\})$.

We will say that the abstract semantics correspond to the concrete semantics of uniform defined similarly as follows.

$$\llbracket \text{uniform } x \ n_1 \ n_2 \rrbracket \delta = (\delta \upharpoonright fv(\delta) - \{x\}) \times \delta_2$$

where $\delta_2 = (\lambda \sigma. \mathbf{if } n_1 \leq \sigma(x) \leq n_2 \mathbf{ then } \frac{1}{n_2 - n_1 + 1} \mathbf{ else } 0)$.

The soundness of the abstract semantics follows immediately from the soundness of forget and product.

²See Appendix D for a precise definition of invertibility.

5) *Conditioning*: Distribution conditioning for probabilistic polyhedra serves the same role as meet in the classic domain of polyhedra in that each is used to perform abstract evaluation of a conditional expression in its respective domain.

Definition 14. Consider the probabilistic polyhedron P_1 and Boolean expression B . Let n, \bar{n} be such that $n = C_1 \# B$ and $\bar{n} = C_1 \# (\neg B)$. The value n is an over-approximation of the number of points in C_1 that satisfy the condition B and \bar{n} is an over-approximation of the number of points in C_1 that do not satisfy B . Then $P_1 \mid B$ is the probabilistic polyhedron P_2 defined as follows.

$$\begin{aligned} p_2^{\min} &= p_1^{\min} & s_2^{\min} &= \max \{s_1^{\min} - \bar{n}, 0\} \\ p_2^{\max} &= p_1^{\max} & s_2^{\max} &= \min \{s_1^{\max}, n\} \\ m_2^{\min} &= \max \{p_2^{\min} \cdot s_2^{\min}, m_1^{\min} - p_1^{\max} \cdot \min \{s_1^{\max}, \bar{n}\}\} \\ m_2^{\max} &= \min \{p_2^{\max} \cdot s_2^{\max}, m_1^{\max} - p_1^{\min} \cdot \max \{s_1^{\min} - n, 0\}\} \\ C_2 &= \langle\langle B \rangle\rangle C_1 \end{aligned}$$

The maximal and minimal probability per point are unchanged, as conditioning simply retains points from the original distribution. To compute the minimal number of points in P_2 , we assume that as many points as possible from C_1 fall in the region satisfying $\neg B$. The maximal number of points is obtained by assuming that a maximal number of points fall within the region satisfying B .

The total mass calculations are more complicated. There are two possible approaches to computing m_2^{\min} and m_2^{\max} . The bound m_2^{\min} can never be less than $p_2^{\min} \cdot s_2^{\min}$, and so we can always safely choose this as the value of m_2^{\min} . Similarly, we can always choose $p_2^{\max} \cdot s_2^{\max}$ as the value of m_2^{\max} . However, if m_1^{\min} and m_1^{\max} give good bounds on total mass (i.e., m_1^{\min} is much higher than $p_1^{\min} \cdot s_1^{\min}$ and dually for m_1^{\max}), then it can be advantageous to reason starting from these bounds.

We can obtain a sound value for m_2^{\min} by considering the case where a maximal amount of mass from C_1 fails to satisfy B . To do this, we compute $\bar{n} = C_1 \# \neg B$, which provides an over-approximation of the number of points within C_1 but outside the area satisfying B . We bound \bar{n} by s_1^{\max} and then assign each of these points maximal mass p_1^{\max} , and subtract this from m_1^{\min} , the previous lower bound on total mass.

By similar reasoning, we can compute m_2^{\max} by assuming a minimal amount of mass m is removed by conditioning, and subtracting m from m_1^{\max} . This m is given by considering an under-approximation of the number of points falling outside the area of overlap between C_1 and B and assigning each point minimal mass as given by p_1^{\min} . This m is given by $\max \{s_1^{\min} - n, 0\}$.

Figure 5 demonstrates the components that affect the conditioning operation. The figure depicts the integer-valued points present in two polyhedra—one representing C_1 and the other representing B (shaded). As the set of points in C_1




Figure 5. Example of distribution conditioning in the abstract domain \mathbb{P} .

satisfying B is convex, this region is precisely represented by $\langle\langle B \rangle\rangle C_1$. By contrast, the set of points in C_1 that satisfy $\neg B$ is not convex, and thus $\langle\langle \neg B \rangle\rangle C_1$ is an over-approximation. The icons beside the main image indicate which shapes correspond to which components and the numbers within the icons give the total count of points within those shapes.

Suppose the components of P_1 are as follows.

$$\begin{aligned} s_1^{\min} &= 19 & p_1^{\min} &= 0.01 & m_1^{\min} &= 0.85 \\ s_1^{\max} &= 20 & p_1^{\max} &= 0.05 & m_1^{\max} &= 0.9 \end{aligned}$$

Then $n = 4$ and $\bar{n} = 16$. Note that we have set \bar{n} to be the number of points in the non-shaded region of Figure 5. This is more precise than the count given by $\#(\langle\langle B \rangle\rangle C)$, which would yield 18. This demonstrates why it is worthwhile to have a separate operation for counting points satisfying a boolean expression. These values of n and \bar{n} give us the following for the first four numeric components of P_2 .

$$\begin{aligned} s_2^{\min} &= \max(19 - 16, 0) = 3 & p_2^{\min} &= 0.01 \\ s_2^{\max} &= \min(20, 4) = 4 & p_2^{\max} &= 0.05 \end{aligned}$$

For the m_2^{\min} and m_2^{\max} , we have the following for the method of calculation based on $p_2^{\min/\max}$ and $s_2^{\min/\max}$.

$$m_2^{\min} = 0.01 \cdot 3 = 0.03 \quad m_2^{\max} = 0.05 \cdot 4 = 0.2$$

For the method of computation based on $m_1^{\min/\max}$, we have

$$\begin{aligned} m_2^{\min} &= 0.85 - 0.05 \cdot 16 = 0.05 \\ m_2^{\max} &= 0.9 - 0.01 \cdot (19 - 4) = 0.75 \end{aligned}$$

In this case, the calculation based on subtracting from total mass provides a tighter estimate for m_2^{\min} , while the method based on multiplying p_2^{\max} and s_2^{\max} is better for m_2^{\max} .

Lemma 15. If $\delta \in \gamma_{\mathbb{P}}(P)$ then $\delta \mid B \in \gamma_{\mathbb{P}}(P \mid B)$.

6) *Scalar Product*: The scalar product is straightforward, as it just scales the mass per point and total mass.

Definition 16. Given a scalar p in $[0, 1]$, we write $p \cdot P_1$ for the probabilistic polyhedron P_2 specified below.

$$\begin{aligned} s_2^{\min} &= s_1^{\min} & p_2^{\min} &= p \cdot p_1^{\min} \\ s_2^{\max} &= s_1^{\max} & p_2^{\max} &= p \cdot p_1^{\max} \\ m_2^{\min} &= p \cdot m_1^{\min} & C_2 &= C_1 \\ m_2^{\max} &= p \cdot m_1^{\max} & & \end{aligned}$$

Lemma 17. If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ then $p \cdot \delta_1 \in \gamma_{\mathbb{P}}(p \cdot P_1)$.

7) *Normalization*: If a probabilistic polyhedron P has $m^{\min} = 1$ and $m^{\max} = 1$ then it represents a normalized distribution. We define below an abstract counterpart to distribution normalization, capable of transforming an arbitrary probabilistic polyhedron into one containing only normalized distributions.

Definition 18. Whenever $m_1^{\min} > 0$, we write $\text{normal}(P_1)$ for the probabilistic polyhedron P_2 specified below.

$$\begin{array}{l|l} p_2^{\min} = p_1^{\min}/m_1^{\max} & s_2^{\min} = s_1^{\min} \\ p_2^{\max} = p_1^{\max}/m_1^{\min} & s_2^{\max} = s_1^{\max} \\ m_2^{\min} = m_2^{\max} = 1 & C_2 = C_1 \end{array}$$

When $m_1^{\min} = 0$, we set $p_2^{\max} = 1$. Note that if P_1 is the zero distribution then $\text{normal}(P_1)$ is not defined.

Lemma 19. If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\text{normal}(\delta_1)$ is defined, then $\text{normal}(\delta_1) \in \gamma_{\mathbb{P}}(\text{normal}(P_1))$.

D. Policy Evaluation

Here we show how to implement the threshold test given as Definition 3 using probabilistic polyhedra. To make the definition simpler, let us first introduce a bit of notation.

Notation 20. If P is a probabilistic polyhedron over variables V , and σ is a state over variables $V' \subseteq V$, then $P \mid \sigma \stackrel{\text{def}}{=} P \mid B$ where $B = \bigwedge_{x \in V'} x = \sigma(x)$.

Definition 21. Given some probabilistic polyhedron P_1 and statement S where $\langle\langle S \rangle\rangle P_1$ terminates, let $P_2 = \langle\langle S \rangle\rangle P_1$ and $P_3 = P_2 \upharpoonright L$. If, for every $\sigma_L \in \gamma_{\mathbb{C}}(C_3)$ with $\neg \text{iszero}(P_2 \mid \sigma_L)$, we have $P_4 = \text{normal}((P_2 \mid \sigma_L) \upharpoonright H)$ with $p_4^{\max} \leq t$, then we write $t\text{secure}_t(S, P_1)$.

The computation of P_3 involves only abstract interpretation and projection, which are computable using the operations defined previously in this section. If we have a small number of outputs (as for the binary outputs considered in our examples), we can enumerate them and check $\neg \text{iszero}(P_2 \mid \sigma_L)$ for each output σ_L . When this holds (that is, the output is feasible), we compute P_4 , which again simply involves the abstract operations defined previously. The final threshold check is then performed by comparing p_4^{\max} to the probability threshold t .

Now we state the main soundness theorem for abstract interpretation using probabilistic polyhedra. This theorem states that the abstract interpretation just described can be used to soundly determine whether to accept a query.

Theorem 22. Let δ be an attacker's initial belief. If $\delta \in \gamma_{\mathbb{P}}(P_1)$ and $t\text{secure}_t(S, P_1)$, then S is threshold secure for threshold t when evaluated with initial belief δ .

VI. POWERSSET OF PROBABILISTIC POLYHEDRA

This section presents the $\mathcal{P}_n(\mathbb{P})$ domain, an extension of the \mathbb{P} domain that abstractly represents a set of distributions as at most n probabilistic polyhedra, elements of \mathbb{P} .

Definition 23. A *probabilistic (polyhedral) set* Δ is a set of probabilistic polyhedra, or $\{P_i\}$ with each P_i over the same variables. We write $\mathcal{P}_n(\mathbb{P})$ for the domain of probabilistic polyhedral powersets composed of no more than n probabilistic polyhedra.

Each probabilistic polyhedron P is interpreted disjunctively: it characterizes one of many possible distributions. The probabilistic polyhedral set is interpreted additively. To define this idea precisely, we first define a lifting of $+$ to sets of distributions. Let D_1, D_2 be two sets of distributions. We then define addition as follows.

$$D_1 + D_2 = \{\delta_1 + \delta_2 \mid \delta_1 \in D_1 \wedge \delta_2 \in D_2\}$$

This operation is commutative and associative and thus we can use \sum for summations without ambiguity as to order of operations. The concretization function for $\mathcal{P}_n(\mathbb{P})$ is then defined as:

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) \stackrel{\text{def}}{=} \sum_{P \in \Delta} \gamma_{\mathbb{P}}(P)$$

We can characterize the condition of Δ containing only the zero distribution, written $\text{iszero}(\Delta)$, via the condition that all of the member probabilistic polyhedra are zero.

$$\text{iszero}(\Delta) \stackrel{\text{def}}{=} \bigwedge_{P \in \Delta} \text{iszero}(P)$$

A. Abstract Semantics for $\mathcal{P}_n(\mathbb{P})$

With a few exceptions, the abstract implementations of the basic operations for the powerset domain are extensions of operations defined on the base probabilistic polyhedra domain.

Theorem 24. For all δ, S, Δ , if $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $\langle\langle S \rangle\rangle \Delta$ terminates, then $\llbracket S \rrbracket \delta$ terminates and $\llbracket S \rrbracket \delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\langle\langle S \rangle\rangle \Delta)$.

Proof of this theorem is given in Appendix E.

Definition 25. The *powerset simplification* transforms a set containing potentially more than n elements into one containing no more than n , for $n \geq 1$. The simplest approach involves repeated use of abstract plus in the base domain \mathbb{P} .

$$\lfloor \{P_i\}_{i=1}^m \rfloor_n \stackrel{\text{def}}{=} \begin{cases} \{P_i\}_{i=1}^m & \text{if } m \leq n \\ \lfloor \{P_i\}_{i=1}^{m-2} \cup \{P_{m-1} + P_m\} \rfloor_n & \text{otherwise} \end{cases}$$

Lemma 26. $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\lfloor \Delta \rfloor_m)$ where $m \leq n$.

Note that the order in which individual probabilistic polyhedra are simplified has no effect on soundness but may impact the precision of the resulting abstraction.

Many of the operations and lemmas for the powerset domain are simple liftings of the corresponding operations and lemmas for single probabilistic polyhedra. For these operations (operations 1-5 given below), we simply list the definition.

1) *Forget*: $f_y(\Delta) \stackrel{\text{def}}{=} \{f_y(P) \mid P \in \Delta\}$

- 2) *Project*: $\Delta \uparrow V \stackrel{\text{def}}{=} \{P \uparrow V \mid P \in \Delta\}$
- 3) *Conditioning*: $\Delta \mid B \stackrel{\text{def}}{=} \{P \mid B \mid P \in \Delta\}$
- 4) *Assignment*: $\Delta [x \rightarrow E] \stackrel{\text{def}}{=} \{P[x \rightarrow E] \mid P \in \Delta\}$
- 5) *Scalar product*: $p \cdot \Delta \stackrel{\text{def}}{=} \{p \cdot P \mid P \in \Delta\}$

6) *Product*: The product operation is only required for the special uniform statement and only applies to the product of a probabilistic set with a single probabilistic polyhedron. $\Delta \times P' \stackrel{\text{def}}{=} \{P \times P' \mid P \in \Delta\}$ (where we assume that $\text{fv}(\Delta) \cap P' = \emptyset$).

7) *Plus*: The abstract plus operation involves simplifying the combined contributions from two sets into one bounded set: $\Delta_1 + \Delta_2 \stackrel{\text{def}}{=} [\Delta_1 \cup \Delta_2]_n$, whenever $\neg \text{iszero}(\Delta_1)$ and $\neg \text{iszero}(\Delta_2)$. Alternatively, if $\text{iszero}(\Delta_1)$ (or $\text{iszero}(\Delta_2)$) then $\Delta_1 + \Delta_2$ is defined to be identical to Δ_2 (or Δ_1).

8) *Normalization*: Since in the $\mathcal{P}_n(\mathbb{P})$ domain, the over(under) approximation of the total mass is not contained in any single probabilistic polyhedron, the normalization must scale each component of a set by the overall total. The minimum (maximum) mass of a probabilistic polyhedra set $\Delta = \{P_1, \dots, P_n\}$ is defined as follows.

$$M^{\min}(\Delta) \stackrel{\text{def}}{=} \sum_{i=1}^n m_i^{\min} \quad | \quad M^{\max}(\Delta) \stackrel{\text{def}}{=} \sum_{i=1}^n m_i^{\max}$$

Definition 27. The scaling of a probabilistic polyhedra P_1 by minimal total mass \underline{m} and maximal total mass \overline{m} , written $\text{normal}(P)(\underline{m}, \overline{m})$ is the probabilistic polyhedron P_2 defined as follows whenever $\underline{m} > 0$.

$$\begin{array}{l|l} p_2^{\min} = p_1^{\min} / \overline{m} & s_2^{\min} = s_1^{\min} \\ p_2^{\max} = p_1^{\max} / \underline{m} & s_2^{\max} = s_1^{\max} \\ m_2^{\min} = m_1^{\min} / \overline{m} & C_2 = C_1 \\ m_2^{\max} = m_1^{\max} / \underline{m} & \end{array}$$

Whenever $\underline{m} = 0$ the resulting P_2 is defined as above but with $p_2^{\max} = 1$ and $m_2^{\max} = 1$.

Normalizing a set of probabilistic polyhedra can be defined as follows

$$\text{normal}(\Delta) \stackrel{\text{def}}{=} \{\text{normal}(P)(M^{\min}(\Delta), M^{\max}(\Delta)) \mid P \in \Delta\}$$

B. Policy Evaluation

Determining the bound on the probability of any state represented by a single probabilistic polyhedron is as simple as checking the p^{\max} value in the normalized version of the probabilistic polyhedron. In the domain of probabilistic polyhedron sets, however, the situation is more complex, as polyhedra may overlap and thus a state's probability could involve multiple probabilistic polyhedra. A simple estimate of the bound can be computed by abstractly adding all the probabilistic polyhedra in the set, and using the p^{\max} value of the result.

Lemma 28. If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $P_1 = \sum_{P \in \Delta} P$ then $\max_{\sigma} \delta(\sigma) \leq p_1^{\max}$.

This approach has an unfortunate tendency to increase the probability bound determined as one increases the bound

on the number of probabilistic polyhedra allowed. A more complicated method, which is used in our implementation, computes a partition of the polyhedra in the set into another set of disjoint polyhedra and determines the maximum probable point among the representatives of each region in the partition. In order to present this method precisely we begin with some definitions.

Definition 29. The (maximum) probability of a state σ according to a probabilistic polyhedron P_1 , written $P_1^{\max}(\sigma)$, is p_1^{\max} if $P_1 \in \gamma_{\mathcal{C}}(C_1)$ and 0 otherwise.

$$P_1^{\max}(\sigma) = \begin{cases} p_1^{\max} & \text{if } \sigma \in \gamma_{\mathcal{C}}(C_1) \\ 0 & \text{otherwise} \end{cases}$$

Likewise the (maximum) probability of σ according to a probabilistic polyhedra set $\Delta = \{P_i\}$, written $\Delta^{\max}(\sigma)$, is defined as follows.

$$\Delta^{\max}(\sigma) = \sum_i P_i^{\max}(\sigma)$$

A mere application of the various definitions allows one to conclude the following remark.

Remark 30. If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ then $\delta(\sigma) \leq \Delta^{\max}(\sigma)$ and therefore $\max_{\sigma} \delta(\sigma) \leq \max_{\sigma} \Delta^{\max}(\sigma)$ for every σ .

Taking advantage of the domain, we will produce a set of representative points $\{\sigma_i\}$ with $\max_i \Delta^{\max}(\sigma_i) = \max_{\sigma} \Delta^{\max}(\sigma)$. To do this, we first need to define a linear partition.

Definition 31. A *poly partition* of a set of polyhedra $\{P_i\}$ is another set of polyhedra $\{L_i\}$, usually of larger size, with the following properties.

- 1) $\gamma_{\mathcal{C}}(L_i) \cap \gamma_{\mathcal{C}}(L_j) = \emptyset$ for every $i \neq j$.
- 2) $\cup_i \gamma_{\mathcal{C}}(L_i) = \cup_i \gamma_{\mathcal{C}}(P_i)$
- 3) For every i, j , either $\gamma_{\mathcal{C}}(L_i) \subseteq \gamma_{\mathcal{C}}(P_j)$ or $\gamma_{\mathcal{C}}(L_i) \cap \gamma_{\mathcal{C}}(P_j) = \emptyset$.

Any set $\{\sigma_i\}$, with $\sigma_i \in \gamma_{\mathcal{C}}(L_i)$ for every i , will be called a *representative set* of the partition.

We can now determine the maximal probability using only representative points, one from each piece of the poly partition.

Lemma 32. $\max_{\sigma \in R} \Delta^{\max}(\sigma) = \max_{\sigma} \Delta^{\max}(\sigma)$ where \mathcal{L} is a poly partition of Δ and R is a representative set of \mathcal{L} .

Note that the set of representatives R is not unique and the lemma holds for any such set.

We will write $\text{max}_{pp}(\Delta)$ for $\max_{\sigma} \Delta^{\max}(\sigma)$ to make explicit the method with which this value can be computed according to the lemma above.

Notation 33. If Δ is a probabilistic polyhedron set over variables V , and σ is a state over variables $V' \subseteq V$, then $\Delta \mid \sigma \stackrel{\text{def}}{=} \Delta \mid B$ where $B = \bigwedge_{x \in V'} x = \sigma(x)$.

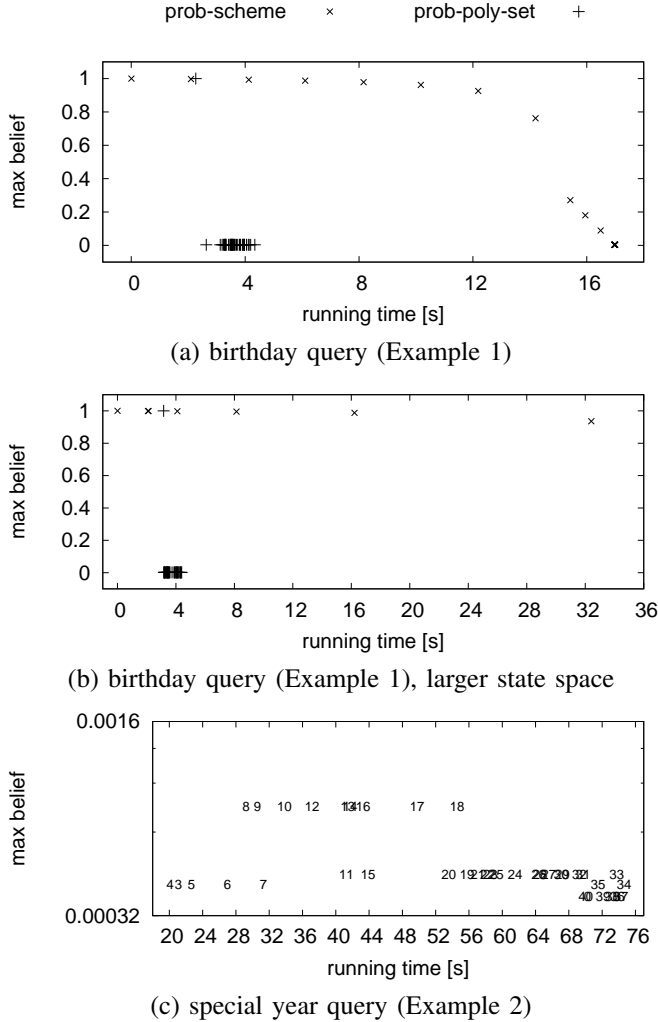


Figure 6. Query evaluation comparison

Definition 34. Given some probabilistic polyhedron Δ_1 and statement S where $\langle\langle S \rangle\rangle \Delta_1$ terminates, let $\Delta_2 = \langle\langle S \rangle\rangle \Delta_1$ and $\Delta_3 = \Delta_2 \upharpoonright L = \{P'_i\}$. If for every $\sigma_L \in \gamma_{\mathcal{P}(\mathbb{C})}(\{C'_i\})$ with $\neg \text{iszero}(\Delta_2 \upharpoonright \sigma_L)$ we have $\Delta_4 = \text{normal}((\Delta_2 \upharpoonright \sigma_L) \upharpoonright H)$ and $\max_{pp}(\Delta_4) \leq t$, then we write $t\text{secure}_t(S, \Delta_1)$.

Below we state the main soundness theorem for abstract interpretation using probabilistic polyhedron sets. This theorem states that the abstract interpretation just described can be used to soundly determine whether to accept a query.

Theorem 35. Let δ be an attacker’s initial belief. If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $t\text{secure}_t(S, \Delta)$, then S is threshold secure for threshold t when evaluated with initial belief δ .

VII. IMPLEMENTATION AND EXPERIMENTS

We have implemented an interpreter for the core language based on the probabilistic polyhedra powerset domain. The base manipulations of polyhedra are done using the Parma

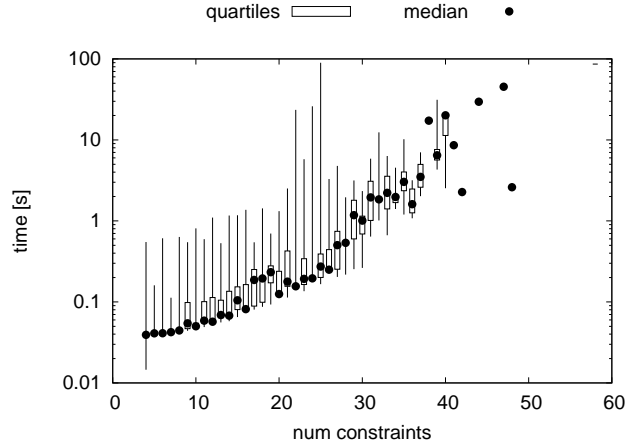


Figure 7. LattE benchmarks

Polyhedra Library [17]. Size calculations are done using the LattE lattice point counter [18]. LattE is also used for the integer linear programming problem involved in the abstract forget operation. The interpreter itself is written in OCaml. We conducted several experiments on a Mac Pro with two 2.26 GHz quad-core Xeon processors using 16 GB of RAM and running OS X v10.6.7. While many of the abstract operations distribute over the set of probabilistic polyhedra and thus could be parallelized, our implementation is currently single-threaded.

Figure 6(a) illustrates the result of running the query given in Example 1 (Section II) using our implementation and one using Probabilistic Scheme [11], which is capable of sound probability estimation after partial enumeration. Each \times plots prob-scheme’s maximum probability value (the y axis)—that is, the probability it assigns to the most likely secret state—when given a varying amount of time for sampling (the x axis). We can see the precision improves steadily until it reaches the exact value of $1/259$ at around 17 seconds. Each $+$ plots our implementation’s maximum probability value when given an increasing number of probabilistic polyhedra; with a polyhedral bound of 2 (or more), we obtain the exact value in less than 3 seconds. The timing measurements are taken to be the medians of 12 runs. The advantage of our approach is more evident in Figure 6(b) where we use the same program but allow *byear* to span 1910 to 2010 rather than 1956 to 1992. In this case prob-scheme makes little progress even after a minute, and eventually runs out of memory. Our approach, however, is unaffected by this larger state space and produces the exact maximum belief in around 3 seconds when using only 2 probabilistic polyhedra.

Figure 6(c) shows the result of our implementation assessing the special query (Example 2) with initial belief matching that following the first birthday query. Each plotted point is the number of polyhedra allowed. The result

time [s] SIQR (outliers) max belief	prob-poly set size bound																
	1	2	3	4	5	6	7	8	9	10	15	20	25	30	35	40	∞
bday 1	2.0 0.6 (1) 1	2.4 0.6 (1) 0.00386	3.2 0.5 (1) 0.00386	3.2 1.1 (0) 0.00386	3.1 1.0 (0) 0.00386	3.4 0.8 (0) 0.00386	3.2 0.6 (2) 0.00386	3.4 0.7 (0) 0.00386	3.4 0.9 (0) 0.00386	3.7 0.8 (0) 0.00386	3.4 1.1 (0) 0.00386	3.4 1.2 (0) 0.00386	3.3 0.6 (3) 0.00386	3.6 0.9 (0) 0.00386	3.3 0.3 (4) 0.00386	3.1 0.8 (0) 0.00386	3.5 1.1 (0) 0.00386
bday 1+2	4.0 1.1 (0) 1	5.1 0.9 (0) 1	6.8 1.4 (0) 0.02703	6.3 1.8 (0) 0.02703	5.2 2.1 (0) 0.02703	7.0 2.1 (0) 0.02703	6.1 1.4 (0) 0.02703	7.4 1.2 (0) 0.02703	6.8 1.2 (0) 0.02703	7.9 1.3 (0) 0.02703	7.5 1.4 (0) 0.02703	7.0 1.2 (0) 0.02703	6.7 1.6 (0) 0.02703	8.2 1.4 (0) 0.02703	7.3 0.6 (3) 0.02703	6.7 1.7 (0) 0.02703	7.9 1.8 (0) 0.02703
bday 1+2+special	10.9 3.8 (0) 1	14.8 1.3 (3) 1	20.2 2.6 (3) 4.22e-4	18.8 4.9 (0) 4.22e-4	17.7 6.5 (0) 4.22e-4	26.3 7.7 (0) 4.22e-4	30.2 5.5 (0) 4.22e-4	28.3 4.1 (0) 8.06e-4	29.0 4.3 (0) 8.06e-4	32.7 4.9 (0) 8.06e-4	42.5 2.1 (4) 4.60e-4	51.4 2.5 (3) 4.60e-4	57.3 6.4 (2) 4.60e-4	65.4 2.1 (2) 4.60e-4	70.3 2.3 (3) 4.22e-4	67.3 8.8 (1) 3.84e-4	68.7 8.5 (2) 3.84e-4
bday large 1	2.6 1.2 (0) 1	3.2 0.8 (0) 0.00141	3.8 0.7 (1) 0.00141	3.3 0.9 (0) 0.00141	4.3 0.8 (0) 0.00141	3.3 1.2 (0) 0.00141	3.5 0.3 (5) 0.00141	3.3 0.4 (2) 0.00141	3.1 0.4 (5) 0.00141	3.5 1.1 (0) 0.00141	4.1 1.1 (0) 0.00141	3.5 1.1 (0) 0.00141	2.9 0.6 (4) 0.00141	3.3 1.0 (0) 0.00141	2.3 0.4 (2) 0.00141	3.9 0.9 (0) 0.00141	3.9 0.6 (0) 0.00141
bday large 1+2	5.0 1.6 (0) 1	7.2 1.5 (0) 1	9.1 0.5 (4) 0.00990	7.6 1.5 (0) 0.00990	8.6 0.7 (3) 0.00990	6.7 2.1 (0) 0.00990	6.9 1.0 (2) 0.00990	7.8 1.0 (2) 0.00990	7.6 1.3 (0) 0.00990	6.9 2.0 (0) 0.00990	8.1 1.2 (2) 0.00990	7.4 1.3 (0) 0.00990	5.7 2.1 (0) 0.00990	7.2 0.8 (2) 0.00990	6.6 2.0 (0) 0.00990	7.3 1.0 (0) 0.00990	7.1 1.5 (0) 0.00990
bday large 1+2+special	13.2 2.9 (0) 0.00130	17.8 0.9 (2) 3.57e-4	24.1 6.9 (2) 2.11e-4	24.0 7.4 (2) 2.11e-4	27.6 15.3 (3) 2.11e-4	27.6 12.1 (3) 2.11e-4	32.8 6.1 (2) 2.11e-4	31.6 9.0 (2) 4.03e-4	34.5 7.7 (1) 4.03e-4	35.5 16.8 (2) 4.03e-4	49.0 7.0 (2) 2.30e-4	53.3 4.1 (2) 4.03e-4	60.0 7.4 (2) 2.30e-4	72.4 13.8 (3) 2.30e-4	79.4 16.9 (4) 2.30e-4	84.5 6.2 (2) 2.30e-4	87.8 11.2 (3) 2.30e-4
pizza	387.4 26.5 (2) 1	191.0 14.5 (2) 1	161.8 6.9 (2) 1	160.2 7.4 (2) 8.66e-10	311.3 15.3 (3) 4.95e-10	200.9 12.1 (3) 1.50e-10	149.6 6.1 (2) 1.50e-10	151.7 9.0 (2) 1.37e-10	608.0 7.7 (1) 1.37e-10	166.2 16.8 (2) 1.37e-10	221.5 7.0 (2) 6.00e-11	219.8 4.1 (2) 6.00e-11	257.1 7.4 (2) 6.00e-11	306.2 13.8 (3) 6.00e-11	345.1 16.9 (4) 6.00e-11	391.3 6.2 (2) 6.00e-11	374.9 11.2 (3) 6.00e-11
photo	6.8 1.8 (0) 1	4.8 1.0 (0) 0.14286	6.6 2.1 (0) 0.14286	7.7 2.4 (0) 0.14286	7.9 0.9 (1) 0.14286	10.3 1.3 (2) 0.14286	9.3 1.3 (1) 0.14286	7.9 2.2 (0) 0.14286	11.0 1.6 (2) 0.14286	8.8 3.2 (0) 0.14286	11.4 1.3 (3) 0.14286	8.3 3.3 (0) 0.14286	10.2 1.8 (0) 0.14286	10.7 0.9 (2) 0.14286	8.7 1.5 (0) 0.14286	10.4 0.6 (1) 0.14286	11.2 1.1 (1) 0.14286
travel	214.8 7.1 (0) 1	21.4 0.8 (3) 1	34.9 4.1 (2) 1	33.7 6.9 (0) 1	59.6 8.8 (0) 1	46.8 6.6 (0) 0.01111	74.8 1.8 (4) 0.01111	62.0 3.6 (3) 0.00556	77.5 11.4 (0) 0.00556	72.9 5.7 (2) 0.00139	139.2 5.9 (3) 0.00123	133.2 8.0 (3) 0.00101	149.4 22.5 (0) 5.05e-4	163.3 14.0 (2) 5.05e-4	170.2 6.3 (3) 5.05e-4	163.0 15.1 (2) 5.05e-4	235.5 28.1 (2) 5.05e-4

Table I
QUERY EVALUATION BENCHMARKS

demonstrates that more complex queries, specifically ones with many disjunctions in their conditionals, not only slow our approach, but also reduce the precision of the maximum probability value. The example requires 36 polyhedra for exact calculations though as little as 3 produce probabilities near exact. Note that the precision does not increase monotonically with the number of polyhedra—in some cases more polyhedra leads to a less precise result. We conjecture that the occasional worsening of the precision with increase in the number of allowable polyhedra is due to an overly simple means of deciding which polyhedra to merge when performing abstract simplification; we plan to investigate this issue in future work.

Table I tabulates details for the example programs along three other queries we developed based on advertising scenarios; these queries are described in the Appendix C. In each box is the wall clock time for processing (median of 12 runs), the running time’s semi-interquartile range (SIQR), the number of outliers, which are defined to be the points $3 \times$ SIQR below the first quartile or above the third, and the max belief computed (smaller being more accurate). Obvious trends are that running time goes up and max belief goes down as the number of polyhedra increase, by and large. There are exceptions to running time trend, and most are close to the SIQR and so possibly not statistically significant. The most striking exception is the

running time for poly-size 9 of the “pizza” query. This extreme outlier is due to a single invocation of LattE on the largest set of constraints among all the benchmarks performed in the table. We have no good explanation of how this complex polyhedron arose. The only exceptions to monotonic decrease in max belief are the “special queries”, as already discussed.

Investigating the running time results further, we discovered that for nearly all benchmarks, 95% or more of the running time is spent in the LattE counting tool. The LattE tool exhibits super-exponential running time in terms of the number of constraints (see Figure 7) over the polyhedra that occur when evaluating the various queries in Table I. As such, overall running time is susceptible to the complexity of the polyhedra involved, even when they are few in number. The merging operation, while used to keep the number of probabilistic polyhedra below the required bound, also tends to produce more complex polyhedra. These observations suggest a great deal of performance improvement can be gained by simplifying the polyhedra if they become too complex.

VIII. DISCUSSION AND RELATED WORK

Prior work aimed at controlling access to users’ private data has focused on access control policies. For example, Persona [6] users can store personal data on distributed

storage servers that use attribute-based encryption; only those parties that have the attribute keys for particular data items may see them. Our approach relaxes the access control model to offer more fine-grained information release policies by directly modeling an attacker’s belief.

Others have considered how an adversary’s knowledge of private data might be informed by a program’s output. Clark, Hunt, and Malacaria [27] define a static analysis that bounds the secret information a straight-line program can leak in terms of equivalence relations between the inputs and outputs. Backes et al. [21] automate the synthesis of such equivalence relations and quantify leakage by computing the exact size of equivalence classes. Köpf and Rybalchenko [22] extend this approach, improving its scalability by using sampling to identify equivalence classes and using under- and over-approximation to obtain bounds on their size. Mu and Clark [28] present a similar analysis that uses over-approximation only. In all cases, the inferred equivalence classes can be used to compute entropy-based metrics of information leakage.

We differ from this work in two main ways. First, we implement a different security criterion. The most closely related metric is *vulnerability* V as proposed by Smith [20], which can be defined using our notation as follows:³

Definition 36. Let $\delta' = \llbracket S \rrbracket \delta$, where δ is the model of the querier’s initial belief, and let $\overline{\delta}_X \stackrel{\text{def}}{=} \text{normal}(\delta \upharpoonright X)$. Then query S is *vulnerability threshold secure* iff for

$$V = \sum_{\sigma_L \in \text{support}(\delta'_L)} \overline{\delta}'_L(\sigma_L) \cdot \max_{\sigma_H \in \text{State}_H} \overline{(\delta'|\sigma_L)}_H(\sigma_H)$$

we have $V \leq t$ for some threshold t .

The above definition is an *expectation* over all possible outputs σ_L , so unlikely outputs have less influence. Our notion of threshold security (Definition 3) is stronger because it considers each output individually: if *any* output, however unlikely, would increase knowledge beyond the threshold, the query would be rejected. For example, recall the query from Example 1 where the secret data $bday$ is (assumed by the querier to be) uniformly distributed; call this query Q_1 . According to Definition 36, the minimum acceptable threshold $t \geq V = 2/365 \approx 0.005$, whereas according to Definition 3, the minimum threshold is $t \geq 1/7 \approx 0.143$ which corresponds the equivalence class $260 \leq bday < 267$.

The other main difference is that we keep an on-line model of knowledge according to prior, actual query results, which increases our precision. To see the benefit consider performing query Q_1 followed by a query Q_2 which uses the code from Example 1 but has $today = 265$. With our system and $bday = 270$ the answer to Q_1 is False and with the revised belief the query Q_2 will be accepted as below threshold $t_d = 0.2$. If instead we had to model this

pair of queries statically they would be rejected because (under the assumption of uniformity) the pair of outputs True, True is possible and implies $bday \in \{265, 266\}$ which would require $t_d \geq 0.5$. Our approach also inherits from the belief-based approach the ability to model a querier who is misinformed or incorrect, which can arise following the result of a probabilistic query (more on this below) or because of a change to the secret data between queries [9]. On the other hand, these advantages of our approach come at the cost of maintaining on-line belief models.

Our proposed abstract domains \mathbb{P} and $\mathcal{P}_n(\mathbb{P})$ are useful beyond the application of belief-based threshold security; e.g., they could be used to model uncertainty off-line (as in the above work) rather than beliefs on-line, with the advantage that they are not limited to uniform distributions (as required by [21], [22]). Prior work on probabilistic abstract interpretation is insufficient for this purpose. For example, Monniaux [29] gives an abstract interpretation for probabilistic programs based on over-approximating probabilities. That work contains no treatment of distribution conditioning and normalization, which are crucial for belief-based information flow analysis. The use of under-approximations, needed to soundly handle normalization, is unique to our approach.

McCamant and Ernst’s FLOWCHECK tool [19] measures the information released by a particular execution. However, it measures information release in terms of *channel capacity*, rather than remaining uncertainty which is more appropriate for our setting. For example, FLOWCHECK would report a query that tries to guess a user’s birthday leaks one bit regardless of whether the guess was successful, whereas the belief-based model (and the other models mentioned above) would consider a failing guess to convey very little information (much less than a bit), and a successful guess conveying quite a lot (much more than a bit).

To avoid reasoning directly about an adversary’s knowledge, Dwork and colleagues proposed *differential privacy* [24]: a differentially private query over a database of individuals’ records is a randomized function that produces roughly the same answer whether a particular individual’s data is in the database or not. Thus, if the database curator is trustworthy, there is little reason for an individual to not supply his data. However, we prefer users to control access to their data as they like, rather than have to trust a curator.

In any case, it is difficult to see how to effectively adapt differential privacy, which was conceived for queries over many records, to queries over an individual’s record, as in our setting. To see why, consider the birthday query from Example 1. Bob’s birthday being/not being in the query range influences the output of the query only by 1 (assuming yes/no is 1/0). One could add an appropriate amount of (Laplacian) noise to the query answer to hide what the true answer was and make the query differentially private. However, this noise would be so large compared to the original

³Smith actually proposes *min entropy*, which is $-\log V$.

range $\{0, 1\}$ that the query becomes essentially useless—the user would be receiving a birthday announcement most days.⁴ By contrast, our approach permits answering queries exactly if the release of information is below the threshold. Moreover, there is no limit on the number of queries as long the information release remains bounded; differential privacy, in general, must impose an artificial limit (termed the *privacy budget*) because it does not reason about the information released.

Nevertheless, differential privacy is appealing, and it would be fruitful to consider how to apply its best attributes to our setting. Rastogi and Suciu [23] propose a property called *adversarial privacy* that suggests a way forward. Like our approach, adversarial privacy is defined in terms of a change in attacker knowledge. Roughly: a query’s output on any database may increase an attacker’s a priori belief $\delta(\sigma)$ about any state σ by at most ϵ for all $\delta \in D$ for some $D \in \mathcal{P}(\mathbf{Dist})$. Rastogi and Suciu show that, for a certain class D , adversarial privacy and differential privacy are equivalent, and by relaxing the choice of D one can smoothly trade off utility for privacy. We can take the reverse tack: by modeling a (larger) set of beliefs we can favor privacy over utility. Our abstractions \mathbb{P} and $\mathcal{P}_n(\mathbb{P})$ already model sets of distributions, rather than a single distribution, so it remains interesting future work to exploit this representation toward increasing privacy.

Another important open question for our work is means to handle collusion. Following our motivating example in the Introduction, the user’s privacy would be thwarted if he shared only his birth day with querier X and only his birth year with Y but then X and Y shared their information. A simple approach to preventing this would be to model adversary knowledge globally, effectively assuming that all queriers share their query results; doing so would prevent either X ’s or Y ’s query (whichever was last). This approach is akin to having a global privacy budget in differential privacy and, as there, obviously harms utility. Dealing with collusion is more problematic when using probabilistic queries, e.g., Example 2. This is because highly improbable results make a querier more uncertain, so combining querier knowledge can misrepresent individual queriers’ beliefs. Roughly speaking, querier X could perform a query Q that misinforms the modeled global belief, but since querier Y ’s actual belief is not changed by the result of Q (since he did not actually see its result), he could submit Q' and learn more than allowed by the threshold. Disallowing probabilistic queries solves this problem but harms expressiveness. Another option is to more actively track a set of beliefs, as hinted at above.

⁴By our calculations, with privacy parameter $\epsilon = 0.1$ recommended by Dwork [24], the probability the query returns the correct result is approximately 0.5249.

IX. CONCLUSION

This paper has explored the idea of *knowledge-based security policies*: given a query over some secret data, that query should only be answered if doing so will not increase the querier’s knowledge above a fixed threshold. We enforce knowledge-based policies by explicitly tracking a model of a querier’s belief about secret data, represented as a probability distribution, and we deny any query that could increase knowledge above the threshold. Our denial criterion is independent of the actual secret, so denial does not leak information. We implement query analysis and belief tracking via abstract interpretation using novel domains of *probabilistic polyhedra* and *powersets of probabilistic polyhedra*. Compared to typical approaches to implementing belief revision, our implementation using this domain is more efficient and scales better.

Acknowledgments: The authors are grateful to Jeff Foster, Nataliya Guts, the anonymous referees, and especially Boris Köpf for their helpful comments on drafts of this paper.

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorised to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon. Support was also provided by National Science Foundation grants 0905419 and 0910530.

REFERENCES

- [1] “Facebook developers,” <http://developers.facebook.com/>, 2011, see the policy and docs/guides/canvas directories for privacy information.
- [2] S. Guha, B. Cheng, and P. Francis, “Challenges in measuring online advertising systems,” in *IMC*, 2010.
- [3] “Statement of rights and responsibilities,” <http://www.facebook.com/terms.php>, Oct. 2010.
- [4] D. Worthington, “Myspace user data for sale,” *PC World on-line*, Mar. 2010, http://www.pcworld.com/article/191716/myspace_user_data_for_sale.html.
- [5] S.-W. Seong, J. Seo, M. Nasielski, D. Sengupta, S. Hangal, S. K. Teh, R. Chu, B. Dodson, and M. S. Lam, “PrPI: a decentralized social networking infrastructure,” in *1st International Workshop on Mobile Cloud Computing and Services: Social Networks and Beyond*, Jun. 2010, invited Paper.
- [6] R. Baden, A. Bender, N. Spring, B. Bhattacharjee, and D. Starin, “Persona: an online social network with user-defined privacy,” in *SIGCOMM*, 2009.
- [7] P. Golle, “Revisiting the uniqueness of simple demographics in the us population,” in *WPES*, 2006.

- [8] “Facebook demographics and statistics report 2010,” <http://www.istrategylabs.com/2010/01/facebook-demographics-and-statistics-report-2010-145-growth-in-1-year/>, 2010.
- [9] M. R. Clarkson, A. C. Myers, and F. B. Schneider, “Quantifying information flow with beliefs,” *J. Comput. Secur.*, vol. 17, no. 5, 2009.
- [10] A. Pfeffer, “The design and implementation of IBAL: A general-purpose probabilistic language,” in *Statistical Relational Learning*, L. Getoor and B. Taskar, Eds. MIT Press, 2007.
- [11] A. Radul, “Report on the probabilistic language Scheme,” in *DLS*, 2007.
- [12] S. Park, F. Pfenning, and S. Thrun, “A probabilistic language based on sampling functions,” *TOPLAS*, vol. 31, pp. 4:1–4:46, 2008.
- [13] N. D. Goodman, V. K. Mansinghka, D. M. Roy, K. Bonawitz, and J. B. Tenenbaum, “Church: a language for generative models,” in *UAI*, 2008.
- [14] O. Kiselyov and C.-C. Shan, “Embedded probabilistic programming,” in *DSL*, 2009.
- [15] P. Cousot and N. Halbwachs, “Automatic discovery of linear restraints among variables of a program,” in *POPL*, 1978.
- [16] D. Monniaux, “Abstract interpretation of probabilistic semantics,” in *Static Analysis Symposium (SAS’00)*, no. 1824, 200, pp. 322–339.
- [17] “PPL: The Parma polyhedral library,” <http://www.cs.unipr.it/ppl/>, 2011.
- [18] J. Loera, D. Haws, R. Hemmecke, P. Huggins, J. Tauzer, and R. Yoshida, “Latte,” <http://www.math.ucdavis.edu/latte/>, 2008.
- [19] S. McCamant and M. D. Ernst, “Quantitative information flow as network flow capacity,” in *PLDI*, 2008.
- [20] G. Smith, “On the foundations of quantitative information flow,” in *FOSSACS*, 2009.
- [21] M. Backes, B. Köpf, and A. Rybalchenko, “Automatic discovery and quantification of information leaks,” in *Security and Privacy*, 2009.
- [22] B. Köpf and A. Rybalchenko, “Approximation and randomization for quantitative information-flow analysis,” in *CSF*, 2010.
- [23] V. Rastogi, M. Hay, G. Miklau, and D. Suciu, “Relationship privacy: output perturbation for queries with joins,” in *PODS*, 2009.
- [24] C. Dwork, “A firm foundation for private data analysis,” *Commun. ACM*, vol. 54, pp. 86–95, Jan. 2011.
- [25] P. Mardziel, S. Magill, M. Hicks, and M. Srivatsa, “Dynamic enforcement of knowledge-based security policies,” University of Maryland Department of Computer Science, Tech. Rep. CS-TR-4978, Apr. 2011.
- [26] —, “Dynamic enforcement of knowledge-based security policies,” in *CSF*, Jun. 2011.
- [27] D. Clark, S. Hunt, and P. Malacaria, “Quantitative information flow, relations and polymorphic types,” *J. Log. and Comput.*, vol. 15, pp. 181–199, Apr. 2005.
- [28] C. Mu and D. Clark, “An interval-based abstraction for quantifying information flow,” *Electron. Notes Theor. Comput. Sci.*, vol. 253, pp. 119–141, November 2009.
- [29] D. Monniaux, “Analyse de programmes probabilistes par interprétation abstraite,” Thèse de doctorat, Université Paris IX Dauphine, 2001, résumé étendu en français. Contents in English.

APPENDIX A.

CONCRETE PROBABILISTIC SEMANTICS

Here we briefly explain the concrete probabilistic semantics given in Figure 3. More details can be found in Clarkson et al. [9].

The semantics of `skip` is straightforward: it is the identity on distributions. The semantics of sequences $S_1 ; S_2$ is also straightforward: the distribution that results from executing S_1 with δ is given as input to S_2 to produce the result.

The semantics of assignment is $\delta[x \rightarrow E]$, which is defined as follows:

$$\delta[x \rightarrow E] \stackrel{\text{def}}{=} \lambda\sigma. \sum_{\tau \mid \tau[x \rightarrow \llbracket E \rrbracket \tau] = \sigma} \delta(\tau)$$

In words, the result of substituting an expression E for x is a distribution where state σ is given a probability that is the sum of the probabilities of all states τ that are equal to σ when x is mapped to the distribution on E in τ . For implementation purposes, it will be useful to consider separately the case where assignment is invertible.

When $x \rightarrow E$ is an invertible transformation, the formula for assignment can be simplified to the following, where $x \rightarrow E'$ is the inverse of $x \rightarrow E$.

$$\delta[x \rightarrow E] \stackrel{\text{def}}{=} \lambda\sigma. \delta(\sigma[x \rightarrow \llbracket E' \rrbracket \sigma])$$

When $x \rightarrow E$ is not invertible, the original definition is equivalent to a projection followed by an assignment. Let $V' = \text{domain}(\delta) - \{x\}$ and let $\delta' = \delta \upharpoonright V'$. Then we have the following for a non-invertible assignment.

$$\delta[x \rightarrow E] \stackrel{\text{def}}{=} \lambda\sigma. \text{if } \sigma(x) = \llbracket E \rrbracket \sigma \text{ then } \delta'(\sigma \upharpoonright V') \text{ else } 0$$

In the appendix, we show that this definition by cases is equivalent to the original definition (Theorem 56).

The semantics for conditionals makes use of two operators on distributions which we now define. First, given distributions δ_1 and δ_2 we define the *distribution sum* as follows:

$$\delta_1 + \delta_2 \stackrel{\text{def}}{=} \lambda\sigma. \delta_1(\sigma) + \delta_2(\sigma)$$

In words, the probability mass for a given state σ of the summed distribution is just the sum of the masses from the input distributions for σ . Second, given a distribution δ and a boolean expression B , we define the *distribution conditioned on B* to be

$$\delta|B \stackrel{\text{def}}{=} \lambda\sigma. \text{ if } \llbracket B \rrbracket\sigma \text{ then } \delta(\sigma) \text{ else } 0$$

In short, the resulting distribution retains only the probability mass from δ for states σ in which B holds.

With these two operators, the semantics of conditionals can be stated simply: the resulting distribution is the sum of the distributions of the two branches, where the first branch's distribution is conditioned on B being true, while the second branch's distribution is conditioned on B being false.

The semantics for probabilistic conditionals like that of conditionals but makes use of *distribution scaling*, which is defined as follows: given δ and some scalar p in $[0, 1]$, we have

$$p \cdot \delta \stackrel{\text{def}}{=} \lambda\sigma. p \cdot \delta(\sigma)$$

In short, the probability ascribed to each state is just the probability ascribed to that state by δ but multiplied by p . For probabilistic conditionals, we sum the distributions of the two branches, scaling them according to the odds q and $1 - q$.

The semantics of a single iteration of a while loop is essentially that of if B then S else skip and the semantics of the entire loop is the fixed point of a function that composes the distributions produced by each iteration. That such a fixed point exists is proved by Clarkson et al. [9].

Finally, the semantics of uniform $x \ n_1 \ n_2$, introduced in Section V is given as

$$\llbracket \text{uniform } x \ n_1 \ n_2 \rrbracket\delta = (\delta \upharpoonright V - \{x\}) \times \delta'$$

Where V is the set of variables of δ , and δ' is defined as follows.

$$\delta' = \lambda\sigma. \text{ if } n_1 \leq \sigma(x) \leq n_2 \text{ then } \frac{1}{n_2 - n_1 + 1} \text{ else } 0$$

APPENDIX B.

ALTERNATIVE (FLAWED) THRESHOLD SECURITY POLICY

As an alternative to Definition 3, suppose we used the following instead:

$$\forall \sigma_L \in \text{support}(\delta' \upharpoonright L). (\text{normal}(\delta'|\sigma_L \upharpoonright H))(\sigma_H) \leq t$$

Here is an example that illustrates why this definition is not safe, as it could underestimate the information a querier can learn.

Suppose Bob's threshold for his birth year *byear* is $t = 0.05$. He models a social networking site X as believing

his age is more likely between 20 and 40 than between 40 and 60, e.g., $1971 \leq \text{byear} < 1991$ with probability 0.6 (thus, 0.03 per possibility) and $1951 \leq \text{byear} < 1971$ with probability 0.4 (thus, 0.02 per possibility). If user Bob was born in 1965, then X 's believes his is actual birth year not as likely a more recent year, say 1975; in any case X does not currently believe any possibility above Bob's threshold. Now suppose X submits program S that determines whether Bob's birth year is even. The revised belief will include only even (when *output* = True) or odd (when *output* = False) birthdays, increasing the likelihood of years in the range $[1971, 1991)$ to be 0.06 per point, and the likelihood of years in the range $[1951, 1971)$ to be 0.04 per point. Bob's birthday is 1965, and its probability 0.04 is less than t , so according to the flawed definition the agent would respond to this query. But if this query result is returned, X will see that there are ten possibilities of birth year that are above Bob's threshold. X can deduce that none of these possibilities is Bob's actual birth year, or else the query would have been rejected. Excluding these possibilities, he knows that Bob's birth year is one of ten possibilities between 1951 and 1971 ascribing to each a probability 0.1 which exceeds Bob's threshold of 0.05.

APPENDIX C.

EXAMPLE QUERIES

We provide here the queries and prebeliefs we used for the experiments in Section VII. The queries are described as functions from some set of inputs to some set of outputs. The exact syntax is as follows.

```
querydef queryname in1 ... inn → out1 ... outm :
  querybody
```

To specify a query invocation we use the following syntax.

```
query queryname :
  in1 := val1;
  ...
  inn := valn
```

Each experiment must also specify the values of the secrets being queried, and the querier's prebelief. Each specification is a merely a program that sets the values of these variables. For the actual secret values this program begins with the declaration *secret*; the resulting state of executing program is taken to be the secret state. The program to set the prebelief begins *belief* and has the same format; note that this program will use *pif* or *uniform x n₁ n₂* to give secrets different possible values with different probabilities.

We now give the content of the queries used in the experiments.

1) *Birthday*: For the small stateset size birthday experiments we used the following secret and prebelief.

```
secret :
  s_bday := 270 ;
  s_byear := 1980

belief :
  uniform s_bday 0 364 ;
  uniform s_byear 1956 1992
```

The two queries used were as follows.

```
querydef bday : c_day → output
if s_bday ≥ c_day ∧ c_day + 7 > s_bday then
  output := 1
else
  output := 0
```

```
querydef spec : c_year → output
```

```
age := c_year - s_byear ;
if age = 10 ∨ age = 20 ∨ age = 30 ∨ age = 40 ∨ age = 50 then
  output_temp := 1
else
  output_temp := 0 ;
pif 1/10 then
  output := 1
else
  output := output_temp
```

The statistics shown include the time spent processing this initial setup as well as the following sequences of queries.

- A single bday query alone.

```
query bday :
  c_day := 260
```

- Two bday queries.

```
query bday :
  c_day := 261
```

- Two bday queries followed by a spec query.

```
query spec :
  c_year := 2011
```

2) *Birthday (large)*: For the larger statespace birthday example we used the following secret and prebelief generators.

```
secret :
  s_bday := 270 ;
  s_byear := 1980

belief :
  uniform s_bday 0 364 ;
  uniform s_byear 1910 2010
```

The queries used were identical to the ones for the smaller statespace birthday example.

3) *Pizza*: The pizza example is slightly more complicated, especially in the construction of the prebelief. This example models a targeted Facebook advertisement for a local pizza shop. There are four relevant secret values. The level of school currently being attended by the Facebook user is given by `s_in_school_type`, which is an integer ranging from 0 (not in school) to 6 (Ph.D. program). Birth year is as before and `s_address_lat` and `s_address_long` give the latitude and longitude of the user's home address (represented as decimal degrees scaled by a factor of 10^6 and converted to an integer).

The initial belief models the fact that each subsequent level of education is less likely and also captures the correlation between current educational level and age. For example, a user is given an approximately 0.05 chance of currently being an undergraduate in college, and college attendees are assumed to be born no later than 1985 (whereas elementary school students may be born as late as 2002).

rendering latex

```
secret :
  s_in_school_type := 4 ;
  s_birthday_year := 1983 ;
  s_address_lat := 39003178 ;
  s_address_long := -76958199

belief :
  pif 4/24 then
    uniform s_in_school_type 1 1 ;
    uniform s_birthday_year 1998 2002
  else
    pif 3/19 then
      uniform s_in_school_type 2 2 ;
      uniform s_birthday_year 1990 1998
    else
      pif 2/15 then
        uniform s_in_school_type 3 3 ;
        uniform s_birthday_year 1985 1992
      else
        pif 1/12 then
          uniform s_in_school_type 4 4 ;
          uniform s_birthday_year 1980 1985
        else
          uniform s_in_school_type 0 0 ;
          uniform s_birthday_year 1900 1985 ;
          uniform s_address_lat 38867884 39103178 ;
          uniform s_address_long -77058199 -76825926
```

The query itself targets the pizza advertisement at users who are either 1) in college, 2) aged 18 to 28, or 3) close to the pizza shop (within a square region that is 2.5 miles on each side and centered on the pizza shop). If any of these conditions are satisfied, then the query returns 1, indicating that the ad can be displayed. The full text of the query is given below.

```
querydef pizza : → output
```

```

if  $s\_in\_school\_type \geq 4$  then
   $in\_school := 1$ 
else
   $in\_school := 0$ ;
 $age := 2010 - s\_birth\_year$ ;
if  $age \geq 18 \wedge age \leq 28$  then
   $age\_criteria := 1$ 
else
   $age\_criteria := 0$ ;
 $lr\_lat := 38967884$ ;
 $ul\_lat := 39003178$ ;
 $lr\_long := -76958199$ ;
 $ul\_long := -76925926$ ;
if  $s\_address\_lat \leq ul\_lat \wedge$ 
   $s\_address\_lat \geq lr\_lat \wedge$ 
   $s\_address\_long \geq lr\_long \wedge$ 
   $s\_address\_long \leq ul\_long$  then
   $in\_box := 1$ 
else
   $in\_box := 0$ ;
if  $(in\_school = 1 \vee age\_criteria = 1) \wedge$ 
   $in\_box = 1$  then
   $output := 1$ 
else
   $output := 0$ 

```

4) *Photo*: The photo query is a direct encoding of a case study that Facebook includes on their advertising information page [?]. The advertisement was for CM Photographics, and targets offers for wedding photography packages at women between the ages of 24 and 30 who list in their profiles that they are engaged. The secret state consists of birth year, as before, gender (0 indicates male, 1 indicates female), and “relationship status,” which can take on a value from 0 to 9. Each of these relationship status values indicates one of the status choices permitted by the Facebook software. The example below involves only four of these values, which are given below.

- 0 No answer
- 1 Single
- 2 In a relationship
- 3 Engaged

The secret state and prebelief are as follows.

```

secret :
   $s\_birth\_year := 1983$ ;
   $s\_gender := 0$ ;
   $s\_relationship\_status := 0$ 

belief :
  uniform  $s\_birth\_year$  1900 2010;
  uniform  $s\_gender$  0 1;
  uniform  $s\_relationship\_status$  0 3

```

The query itself is the following.

```

querydef  $cm\_advert : \rightarrow output$ 
 $age := 2010 - s\_birth\_year$ ;
if  $age \geq 24 \wedge age \leq 30$  then
   $age\_sat := 1$ 
else
   $age\_sat := 0$ ;
if  $s\_gender = 1 \wedge$ 
   $s\_relationship\_status = 3 \wedge$ 
   $age\_sat = 1$  then
   $output := 1$ 
else
   $output := 0$ 

```

5) *Travel*: This example is another Facebook advertising case study [?]. It is based on an ad campaign run by Britain’s national tourism agency, VisitBritain. The campaign targeted English-speaking Facebook users currently residing in countries with strong ties to the United Kingdom. They further filtered by showing the advertisement only to college graduates who were at least 21 years of age.

We modeled this using four secret values: country, birth year, highest completed education level, and primary language. As with other categorical data, we represent language and country using an enumeration. We ranked countries by number of Facebook users as reported by socialbakers.com. This resulted in the US being country number 1 and the UK being country 3. To populate the list of countries with “strong connections” to the UK, we took a list of former British colonies. For the language attribute, we consider a 50-element enumeration where 0 indicates “no answer” and 1 indicates “English” (other values appear in the prebelief but are not used in the query).

```

secret :
   $country := 1$ ;
   $birth\_year := 1983$ ;
   $completed\_school\_type := 4$ ;
   $language := 5$ 

belief :
  uniform  $country$  1 200;
  uniform  $birth\_year$  1900 2011;
  uniform  $language$  1 50;
  uniform  $completed\_school\_type$  0 5

querydef  $travel : \rightarrow output$ 
if  $country = 1 \vee country = 3 \vee$ 
   $country = 8 \vee country = 10 \vee$ 
   $country = 18$  then
   $main\_country := 1$ 
else
   $main\_country := 0$ ;
if  $country = 169 \vee country = 197 \vee$ 
   $country = 194 \vee country = 170 \vee$ 
   $country = 206 \vee country = 183 \vee$ 
   $country = 188$  then
   $island := 1$ 
else
   $island := 0$ ;
 $age := 2010 - birth\_year$ ;
if  $language = 1 \wedge$ 
   $(main\_country = 1 \vee island = 1) \wedge$ 
   $age \geq 21 \wedge$ 
   $completed\_school\_type \geq 4$  then
   $output := 1$ 
else
   $output := 0$ 

```

APPENDIX D.
SOUNDNESS PROOFS FOR \mathbb{P}

A. Projection

The proof of projection relies heavily on splitting up the support of a distribution into equivalence classes based on the states they project to. We will have $\sigma, \sigma' \in \text{support}(\delta)$ belonging to the same equivalence class iff $\sigma \upharpoonright V = \sigma' \upharpoonright V$. The details are formalized in the following definition.

Definition 37. *Equivalence classes under projection.*

- $[\sigma_V]_\delta^V$ is an equivalence class of elements of $\text{support}(\delta)$ that project to σ_V (when projected to variables V). Formally, $[\sigma_V]_\delta^V \stackrel{\text{def}}{=} \{\sigma \in \text{support}(\delta) \mid \sigma \upharpoonright V = \sigma_V\}$.
- $[\overline{\sigma_V}]_\delta^V$ is a subset of $\text{support}(\delta)$ that project to anything but σ_V or formally $[\overline{\sigma_V}]_\delta^V \stackrel{\text{def}}{=} \{\sigma \in \text{support}(\delta) \mid \sigma \upharpoonright V \neq \sigma_V\}$.
- $[\sigma_V]_C^V$ is a subset of $\gamma_C(C)$ that project to σ_V (when projected to variables V). Formally, $[\sigma_V]_C^V \stackrel{\text{def}}{=} \{\sigma \in \gamma_C(C) \mid \sigma \upharpoonright V = \sigma_V\}$.
- $[\overline{\sigma_V}]_C^V$ is a subset of $\gamma_C(C)$ that project to anything but σ_V or formally $[\overline{\sigma_V}]_C^V \stackrel{\text{def}}{=} \{\sigma \in \gamma_C(C) \mid \sigma \upharpoonright V \neq \sigma_V\}$.

Remark 38. Let $V \subseteq V' \subseteq \text{fv}(\delta)$, $\sigma \in \text{support}(\delta)$, and $\sigma_V, \sigma'_V \in \text{support}(\delta \upharpoonright V)$.

(i) $(\sigma \upharpoonright V') \upharpoonright V = \sigma \upharpoonright V$

The sets of $\left\{ [\sigma_V]_\delta^V \right\}_{\sigma_V \in \text{support}(\delta \upharpoonright V)}$ form a partition of $\text{support}(\delta)$, equivalently the following two claims.

(ii) $\text{support}(\delta) = \bigcup_{\sigma_V \in \text{support}(\delta \upharpoonright V)} [\sigma_V]_\delta^V$

(iii) $[\sigma_V]_\delta^V \cap [\sigma'_V]_\delta^V = \emptyset$ whenever $\sigma_V \neq \sigma'_V$

Likewise, for any $\sigma_V \in \text{support}(\delta \upharpoonright V)$, the sets of $\left\{ [\sigma_{V'}]_\delta^{V'} \right\}_{\sigma_{V'} \in [\sigma_V]_\delta^V \upharpoonright V'}$ form a partition of $[\sigma_V]_\delta^V$, implying also the following claim.

(iv) $[\sigma_V]_\delta^V = \bigcup_{\sigma_{V'} \in [\sigma_V]_\delta^V \upharpoonright V'} [\sigma_{V'}]_\delta^{V'}$

The equivalence classes in terms of the concrete support sets as related to the abstract support sets are expressed in the following manner.

(v) $[\sigma_V]_\delta^V \subseteq [\sigma_V]_C^V$ and $[\overline{\sigma_V}]_\delta^V \subseteq [\overline{\sigma_V}]_C^V$ whenever $\text{support}(\delta) \subseteq \gamma_C(C)$

Finally, the concrete projection operation can be rewritten in terms of the equivalence classes, a fact we will repeatedly use in the proofs to follow without explicitly stating it.

(vi) $\delta \upharpoonright V = \lambda \sigma_V. \sum_{\sigma \in [\sigma_V]_\delta^V} \delta(\sigma)$

Proof: All of these are merely expansions of the various definitions involved. Note that the two parts of Remark 38 (v) are not contradictory as $[\sigma_V]_\delta^V$ and $[\overline{\sigma_V}]_\delta^V$ are not set complements of each other when viewed as subsets

of $\gamma_C(C)$, though they are complements when viewed as subsets of $\text{support}(\delta)$. \blacksquare

Lemma 39 (Conservation of Mass). *If $V \subseteq \text{fv}(\delta)$ then $\|\delta\| = \|\delta \upharpoonright V\|$.*

Proof: Let us consider the terms of the projected mass sum.

$$\begin{aligned} \|\delta \upharpoonright V\| &= \sum_{\sigma_V \in \text{support}(\delta \upharpoonright V)} \left(\sum_{\sigma \in [\sigma_V]_\delta^V} \delta(\sigma) \right) \\ &= \sum_{\sigma \in \text{support}(\delta)} \delta(\sigma) \\ &= \|\delta\| \end{aligned}$$

The terms in the double sum are the same as those in the single sum as all terms of the first are accounted for in the second due to Remark 38 (ii) and none are double counted due to Remark 38 (iii). \blacksquare

Definition 40. *Concrete forget* can be defined in terms of a projection to all but one variable. That is, $f_x(\delta) \stackrel{\text{def}}{=} \delta \upharpoonright \text{fv}(\delta) - \{x\}$. Also, $f_{x_1, \dots, x_n}(\delta) \stackrel{\text{def}}{=} f_{x_2, \dots, x_n}(f_{x_1}(\delta))$.

The correspondence between repeated concrete forget and a projection involving removal of more than one variable will be demonstrated shortly.

Lemma 41 (Order of Projection). *If $V \subseteq V' \subseteq \text{fv}(\delta)$ then $(\delta \upharpoonright V') \upharpoonright V = \delta \upharpoonright V$.*

Proof: Let $\sigma_V \in \delta \upharpoonright V$.

$$((\delta \upharpoonright V') \upharpoonright V)(\sigma_V) = \sum_{\sigma_{V'} \in [\sigma_V]_\delta^{V'} \upharpoonright V'} \left(\sum_{\sigma \in [\sigma_{V'}]_\delta^{V'}} \delta(\sigma) \right) \quad (1)$$

$$= \sum_{\sigma \in \bigcup_{\sigma_{V'} \in [\sigma_V]_\delta^{V'} \upharpoonright V'} [\sigma_{V'}]_\delta^{V'}} \delta(\sigma) \quad (2)$$

$$= \sum_{\sigma \in [\sigma_V]_\delta^V} \delta(\sigma) \quad (3)$$

$$= (\sigma \upharpoonright V)(\sigma_V) \quad (4)$$

The collapse of the double sums on (1) to (2) is due to the correspondence between the terms of the double sum and the single sum due to Remark 38 (ii) and Remark 38 (iii). The equality of the union of equivalence classes, (2) to (3) is due to Remark 38 (iv). \blacksquare

Corollary 42. $\delta \upharpoonright V = f_{x_1, \dots, x_n}(\delta)$ where $\text{fv}(\delta) - V = \{x_1, \dots, x_n\}$.

Proof: Let us show this by induction on the size of $\overline{V} \stackrel{\text{def}}{=} \text{fv}(\delta) - V$. When $|\overline{V}| = 0$ or $|\overline{V}| = 1$, the claim holds vacuously or by definition of concrete forget, respectively.

Let us assume the claim for $|\overline{V}| = m-1 < n$ and consider the case when $|\overline{V}| = m \leq n$.

$$\begin{aligned}
\delta \upharpoonright V &= (\delta \upharpoonright V \cup \{x_1\}) \upharpoonright V && \text{[by Lemma 41]} \\
&= f_{x_1}(\delta) \upharpoonright V \\
&= f_{x_2, \dots, x_m}(f_{x_1}(\delta)) && \text{[by induction]} \\
&= f_{x_1, \dots, x_m}(\delta)
\end{aligned}$$

Thus, by induction, the claim holds for $m = n$. \blacksquare

Remark 43 (Counting Variations). Two simple counting arguments are required for the further proofs.

- (i) If m objects are distributed fully into two bins, with one of the bins having space for no more than a objects, then the other must have at least $m - a$ objects in it.
- (ii) If m objects are to be packed into bins of sizes a_1, \dots, a_n , with $\sum_i a_i \geq m$, the least number of bins that can be used to fit all m objects is greater or equal to $\lceil m/a^* \rceil$ where $a^* \geq \max_i a_i$.

Proof: Part (i) is immediate. For part (ii), consider some optimal set of bins used to pack the m objects. This set of bins would also let one pack m items assuming each bin had space for exactly a^* objects as this is an upper bound on the size of each bin. Thus the space of solutions to the original packing problem is a subset of the space of solutions to the altered packing problem where all bins are increased to fit a^* items. Thus the solution for the original cannot use fewer bins than the optimal solution for the altered problem. For this alternate problem, the minimum number of bins used to pack all m items is exactly $\lceil m/a^* \rceil$ by a generalization of the pigeonhole principle. \blacksquare

Lemma 7 (Soundness of Forget). *If $\delta \in \gamma_{\mathbb{P}}(P)$ then $f_y(\delta) \in \gamma_{\mathbb{P}}(f_y(P))$.*

Proof: Let $\delta \in \gamma_{\mathbb{P}}(P)$, $V = f_V(\delta) - \{y\}$, and $\delta_2 = \delta \upharpoonright V$. By assumption δ has the following properties.

$$\text{support}(\delta) \subseteq \gamma_{\mathbb{C}}(C) \quad (5)$$

$$s^{\min} \leq |\text{support}(\delta)| \leq s^{\max} \quad (6)$$

$$m^{\min} \leq \|\delta\| \leq m^{\max} \quad (7)$$

$$\forall \sigma \in \text{support}(\delta) . p^{\min} \leq \delta(\sigma) \leq p^{\max} \quad (8)$$

Let $P_2 = f_y(P)$. P_2 thus has the following properties.

$$C_2 = f_y(C) \quad (9)$$

$$p_2^{\min} = p^{\min} \cdot \max \{h_y^{\min} - (\#(C) - s^{\min}), 1\} \quad (10)$$

$$p_2^{\max} = p^{\max} \cdot \min \{h_y^{\max}, s^{\max}\} \quad (11)$$

$$s_2^{\min} = \lceil s^{\min}/h_y^{\max} \rceil \quad (12)$$

$$s_2^{\max} = \min \{\#(C_2), s^{\max}\} \quad (13)$$

$$m_2^{\min} = m^{\min} \quad (14)$$

$$m_2^{\max} = m^{\max} \quad (15)$$

The quantities h_y^{\min} and h_y^{\max} are defined to exhibit the following properties.

$$h_y^{\min} \leq \min_{\sigma_V \in \gamma_{\mathbb{C}}(C_2)} \left| [\sigma_V]_C^V \right| \quad (16)$$

$$h_y^{\max} \geq \max_{\sigma_V \in \gamma_{\mathbb{C}}(C_2)} \left| [\sigma_V]_C^V \right| \quad (17)$$

To show that $\delta_2 \in \gamma_{\mathbb{P}}(f_y(P))$ we need to show the following.

$$\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2) \quad (18)$$

$$s_2^{\min} \leq |\text{support}(\delta_2)| \leq s_2^{\max} \quad (19)$$

$$m_2^{\min} \leq \|\delta_2\| \leq m_2^{\max} \quad (20)$$

$$\forall \sigma_V \in \text{support}(\delta_2) . p_2^{\min} \leq \delta_2(\sigma_V) \leq p_2^{\max} \quad (21)$$

Let us show each of these in turn.

Claim (18) – Support. Let $\sigma_V \in \text{support}(\delta_2)$. Thus $\delta_2(\sigma_V) = \sum_{\sigma \in [\sigma_V]_{\delta}^V} \delta(\sigma) > 0$ so there exists $\sigma \in [\sigma_V]_{\delta}^V$ with $\delta(\sigma) > 0$. So $\sigma \in \text{support}(\delta)$. Therefore, by (5), $\sigma \in \gamma_{\mathbb{C}}(C)$, therefore $\sigma_V \in \gamma_{\mathbb{C}}(C_2)$ by definition of polyhedron forget. Therefore $\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2)$. \square

Claim (19) Support points. First let us show the following claim.

$$\max_{\sigma_V \in \text{support}(\delta_2)} \left| [\sigma_V]_{\delta}^V \right| \leq h_y^{\max} \quad (22)$$

By construction of h_y^{\max} , we have $h_y^{\max} \geq \max_{\sigma_V \in \gamma_{\mathbb{C}}(C_2)} \left| [\sigma_V]_C^V \right|$. Now, $\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2)$ by (18). Also for any $\sigma_V \in \text{support}(\delta_2)$, we have $[\sigma_V]_{\delta}^V \subseteq [\sigma_V]_C^V$ by Remark 38 (v). Therefore $\max_{\sigma_V \in \gamma_{\mathbb{C}}(C_2)} \left| [\sigma_V]_C^V \right| \geq \max_{\sigma_V \in \text{support}(\delta_2)} \left| [\sigma_V]_C^V \right| \geq \max_{\sigma_V \in \text{support}(\delta_2)} \left| [\sigma_V]_{\delta}^V \right|$. Thus concluding $h_y^{\max} \geq \max_{\sigma_V \in \text{support}(\delta_2)} \left| [\sigma_V]_{\delta}^V \right|$.

Consider the elements of $\text{support}(\delta)$ as they map via state projection to elements of $\text{support}(\delta_2)$. Let us view the elements of the later as bins, with the elements of the former as objects to pack into the bins. By (22), we know no bin has more than h_y^{\max} objects, thus we can apply Remark 43 to conclude there are at least $\lceil |\text{support}(\delta)|/h_y^{\max} \rceil$ non-empty bins, or in other words, $|\text{support}(\delta_2)| \geq \lceil |\text{support}(\delta)|/h_y^{\max} \rceil$. This is itself at least as large as $\lceil s^{\min}/h_y^{\max} \rceil = s_2^{\min}$ by (6). Therefore $|\text{support}(\delta_2)| \geq s_2^{\min}$.

For the other side of the inequality, note that the number of bins used, or $|\text{support}(\delta_2)|$ cannot exceed $|\text{support}(\delta)| \leq s^{\max}$ itself. It also cannot exceed $|\gamma_{\mathbb{C}}(C_2)| = \#(C_2)$ given (18). Therefore $\text{support}(\delta_2) \leq \min \{\#(C_2), s^{\max}\}$, concluding requirement (19). \square

Claim (20) Mass. This requirement holds trivially due to Lemma 39 and assumption (7). \square

Claim (21) Probability. Let us first show the following claim. ■

$$\min_{\sigma_V \in \text{support}(\delta_2)} \left| [\sigma_V]_\delta^V \right| \geq h_y^{\min} + s^{\min} - \#(C) \quad (23)$$

Let $\sigma_V \in \text{support}(\delta_2)$. Let us consider the size of $[\overline{\sigma_V}]_\delta^V$.

$$\begin{aligned} \left| [\overline{\sigma_V}]_\delta^V \right| &\leq \left| [\overline{\sigma_V}]_C^V \right| && \text{[by Remark 38 (v)]} \\ &= \#(C) - \left| [\sigma_V]_C^V \right| \\ &\leq \#(C) - \min_{\tau_V \in \gamma_C(C_2)} \left| [\tau_V]_C^V \right| \\ &\leq \#(C) - h_y^{\min} && \text{[by (16)]} \end{aligned}$$

Let us view now the elements of $\text{support}(\delta)$ as mapping (via projection) into two bins, $[\sigma_V]_\delta^V$ and $[\overline{\sigma_V}]_\delta^V$. By the argument above, we know the second bin cannot hold more than $\#(C) - h_y^{\min}$ elements, thus, by Remark 43 (i), it must be the case that the first bin contains at least $|\text{support}(\delta)| - (\#(C) - h_y^{\min})$ elements. This itself is no smaller than $s^{\min} - \#(C) + h_y^{\min}$ by (6). Therefore $\left| [\sigma_V]_\delta^V \right| \geq s^{\min} - \#(C) + h_y^{\min}$ and thus claim (23) holds.

Consider now $\sigma_V \in \text{support}(\delta_2)$. By (8) and the concrete projection definition, it must be the case that $\delta_2(\sigma_V) = \sum_{\sigma \in [\sigma_V]_\delta^V} \delta(\sigma) \geq p^{\min}$. Also,

$$\begin{aligned} \delta_2(\sigma_V) &= \sum_{\sigma \in [\sigma_V]_\delta^V} \delta(\sigma) \\ &\geq \sum_{\sigma \in [\sigma_V]_\delta^V} p^{\min} && \text{[by (8)]} \\ &= \left| [\sigma_V]_\delta^V \right| \cdot p^{\min} \\ &\geq (h_y^{\min} + s^{\min} - \#(C)) \cdot p^{\min} && \text{[by (23)]} \end{aligned}$$

Therefore, $\delta_2(\sigma_V) \geq p^{\min} \cdot \min\{1, h_y^{\min} + s^{\min} - \#(C)\} = p_2^{\min}$, concluding one inequality of the last condition.

For the other inequality, let us once more consider a general $\sigma_V \in \text{support}(\delta_2)$.

$$\begin{aligned} \delta_2(\sigma_V) &= \sum_{\sigma \in [\sigma_V]_\delta^V} \delta(\sigma) \\ &\leq \sum_{\sigma \in [\sigma_V]_\delta^V} p^{\max} && \text{[by (8)]} \\ &= \left| [\sigma_V]_\delta^V \right| \cdot p^{\max} \\ &\leq h_y^{\max} \cdot p^{\max} && \text{[by (22)]} \end{aligned}$$

Since $[\sigma_V]_\delta^V \subseteq \text{support}(\delta)$, we have $\left| [\sigma_V]_\delta^V \right| \leq |\text{support}(\delta)| \leq s^{\max}$ (by (6)). Thus we can also bound $\delta_2(\sigma_V)$ by $s^{\max} \cdot p^{\max}$. Therefore, $\delta_2(\sigma_V) \leq p^{\max} \cdot \min\{h_y^{\max}, s^{\max}\}$, completing the last claim. \square

Lemma 44 (Soundness of Projection). *If $\delta \in \gamma_{\mathbb{P}}(P)$ then $\delta \upharpoonright V \in \gamma_{\mathbb{P}}(P \upharpoonright V)$.*

Proof: Let us show this by induction on the size of $\overline{V} \stackrel{\text{def}}{=} f_V(\delta) - V$. When $|\overline{V}| = 0$ there is no projection to be done, when $|\overline{V}| = 1$, the claim holds by Lemma 7. Let us assume the claim holds for $|\overline{V}| = n - 1$ and look at the case where $|\overline{V}| = n$.

Let us write $\overline{V} = \{x_1, \dots, x_n\}$. Thus $\delta \upharpoonright V = f_{x_1, \dots, x_n}(\delta)$ by Corollary 42. By definition of forget, we also have $f_{x_1, \dots, x_n}(\delta) = f_{x_2, \dots, x_n}(f_{x_1}(\delta))$ and $f_{x_1, \dots, x_n}(P) = f_{x_2, \dots, x_n}(f_{x_1}(P))$. By Lemma 7, we know that $f_{x_1}(\delta) \in \gamma_{\mathbb{P}}(f_{x_1}(P))$, therefore, by induction, $\delta \upharpoonright V = f_{x_2, \dots, x_n}(f_{x_1}(\delta)) \in f_{x_2, \dots, x_n}(f_{x_1}(P)) = P \upharpoonright V$. \blacksquare

B. Assignment

We begin with some useful notation.

Notation 45. Let σ be a state, E be an expression, x be a variable, $S \subseteq \mathbf{State}$, $V \subseteq \mathbf{Var}$.

- $\sigma[x \rightarrow E] \stackrel{\text{def}}{=} \sigma[x \rightarrow \llbracket E \rrbracket \sigma]$
- $S[x \rightarrow E] \stackrel{\text{def}}{=} \{\sigma[x \rightarrow E] \mid \sigma \in S\}$.
- $S \upharpoonright V \stackrel{\text{def}}{=} \{\sigma \upharpoonright V \mid \sigma \in S\}$

Definition 46. A state σ is *feasible* for $x \rightarrow E$ iff $\sigma \in \mathbf{State}[x \rightarrow E]$. We will say that σ is merely feasible if the assignment is clear from the context.

Definition 47. $t_{x \rightarrow E}$ is the function from \mathbf{State} to feasible states (for $x \rightarrow E$) defined by $t_{x \rightarrow E} : \sigma \mapsto \sigma[x \rightarrow E]$

Definition 48. The inverted equivalence class for σ under assignment $x \rightarrow E$ is the set of states that map to σ . We define two varieties, one over all possible states and one for just the states in the support of a distribution.

- $\langle \sigma \rangle_{x \rightarrow E} \stackrel{\text{def}}{=} \{\tau \mid \tau[x \rightarrow E] = \sigma\}$
- $\langle \sigma \rangle_\delta^{x \rightarrow E} \stackrel{\text{def}}{=} \{\tau \in \text{support}(\delta) \mid \tau[x \rightarrow E] = \sigma\}$

Note that σ is feasible iff $\langle \sigma \rangle^{x \rightarrow E} \neq \emptyset$.

Definition 49. An assignment $x \rightarrow E$ is invertible iff $t_{x \rightarrow E}$ is invertible. We will denote $t_{x \rightarrow E}^{-1}$ as the inverse of $t_{x \rightarrow E}$, if it is invertible. The invertibility of $t_{x \rightarrow E}$ is characterized by the existence of the inverse, having the property, that for every $\sigma \in \mathbf{State}$, we have $t_{x \rightarrow E}^{-1}(t_{x \rightarrow E}(\sigma)) = \sigma$. Equivalently, for every feasible state σ , $t_{x \rightarrow E}(t_{x \rightarrow E}^{-1}(\sigma)) = \sigma$.

We can also characterize invertibility via inverted equivalence classes. $x \rightarrow E$ is invertible iff for every feasible σ , $|\langle \sigma \rangle^{x \rightarrow E}| = 1$.

We will say E is invertible if the variable is clear from the context.

Note that since $t_{x \rightarrow E}$ only changes the x component of a state, the inverse, $t_{x \rightarrow E}^{-1}$, also only changes the x component, if the inverse exists. This doesn't mean, however, that the inverse can be represented by an assignment of some E'

to x . Furthermore, since our language for expressions lacks division and non-integer constants, no assignment's inverse can be represented by an assignment.

Definition 50. The expression E is *integer linear* iff $E = n_1 \times x_1 + \dots + n_m \times x_m$, where n_i are integer constants, and x_i are variables. We assume that all the variables in a given context are present. We will generally use x_i and n_i to refer to the contents of a integer linear expression.

From now on, we will assume all expressions E are integer linear. Programs containing non-linear expressions are just not handled by our system at this stage and linear expressions not fully specified are equivalent to integer linear expressions with $n_i = 0$ for variables unused in the original expression.

Lemma 51. $x_1 \rightarrow E$ is non-invertible iff $n_1 = 0$. In other words, $x_1 \rightarrow E$ is non-invertible iff E doesn't depend on x_1 .

Proof: (\Rightarrow) Assume otherwise. Thus E is non-invertible but $n_1 \neq 0$. So we have a feasible state σ with $|\langle \sigma \rangle_{x_1 \rightarrow E}| \neq 1$. Since feasible states have non empty inverted equivalence sets, it must be that $|\langle \sigma \rangle_{x_1 \rightarrow E}| \geq 2$. So let $\tau, \tau' \in \langle \sigma \rangle_{x_1 \rightarrow E}$ with $\tau \neq \tau'$. So $\tau[x_1 \rightarrow E] = \tau'[x_1 \rightarrow E] = \sigma$. Since assignment to x_1 doesn't change the state other than in its value of x_1 , τ and τ' can only differ in their value for x_1 .

But since τ and τ' are identical after the assignment, we have,

$$\begin{aligned} \tau[x_1 \rightarrow E](x_1) &= n_1\tau(x_1) + n_2\tau(x_2) + \dots + n_m\tau(x_m) \\ &= n_1\tau(x_1) + n_2\tau'(x_2) + \dots + n_m\tau'(x_m) \\ &= n_1\tau'(x_1) + n_2\tau'(x_2) + \dots + n_m\tau'(x_m) \\ &= \tau'[x_1 \rightarrow E](x_1) \end{aligned}$$

Canceling out the common $\tau'(x_i)$ terms, we have $n_1\tau(x_1) = n_1\tau'(x_1)$ and since $n_1 \neq 0$, we conclude $\tau(x_1) = \tau'(x_1)$, contradicting $\tau \neq \tau'$.

(\Leftarrow) Let σ be a feasible state and let $\tau \in \langle \sigma \rangle_{x_1 \rightarrow E}$. Let $\tau' = \tau[x_1 \rightarrow \tau(x_1) + 1]$. Since E doesn't depend on x_1 , we have $\llbracket E \rrbracket \tau = \llbracket E \rrbracket \tau'$ and therefore $\tau'[x_1 \rightarrow E] = \tau[x_1 \rightarrow E] = \sigma$ and so we have $\tau, \tau' \in \langle \sigma \rangle_{x_1 \rightarrow E}$ with $\tau \neq \tau'$, therefore E is non-invertible. ■

Lemma 52. Assume $x \rightarrow E$ is non-invertible. σ is feasible iff $\sigma[x \rightarrow E] = \sigma$.

Proof: (\Rightarrow) Let σ be feasible. Thus $\sigma = \tau[x \rightarrow E]$ for some $\tau \in \mathbf{State}$. Since E doesn't depend on x by Lemma 51, we have $(\tau[x \rightarrow E])[x \rightarrow E] = \tau[x \rightarrow E] = \sigma$. So $\sigma[x \rightarrow E] = \sigma$.

(\Leftarrow) Assume $\sigma[x \rightarrow E] = \sigma$. Thus $\sigma \in \mathbf{State}[x \rightarrow E]$ by definition. ■

Lemma 53. Assume $x \rightarrow E$ is non-invertible. Let δ be a distribution with $x \in \text{fv}(\delta)$ and let $V = \text{fv}(\delta) - \{x\}$. If σ is feasible, then $\langle \sigma \rangle_{x \rightarrow E}^V = [\sigma \upharpoonright V]_{\delta}^V$.

Proof: Let $\tau \in \langle \sigma \rangle_{x \rightarrow E}^V$. So $\tau[x \rightarrow E] = \sigma$ and $\tau \in \text{support}(\delta)$. But the assignment only changes x , thus $\tau \upharpoonright V = \sigma \upharpoonright V$, therefore $\tau \in [\sigma \upharpoonright V]_{\delta}^V$. Thus $\langle \sigma \rangle_{x \rightarrow E}^V \subseteq [\sigma \upharpoonright V]_{\delta}^V$.

Let $\tau \in [\sigma \upharpoonright V]_{\delta}^V$. So $\tau \in \text{support}(\delta)$ and $\tau \upharpoonright V = \sigma \upharpoonright V$. Since E doesn't depend on x , we have $\tau[x \rightarrow E] = \sigma[x \rightarrow E] = \sigma$; the second equality follows from Lemma 52 as σ is feasible by assumption. So $\tau \in \langle \sigma \rangle_{x \rightarrow E}^V$. Therefore $[\sigma \upharpoonright V]_{\delta}^V \subseteq \langle \sigma \rangle_{x \rightarrow E}^V$. ■

Remark 54. Assume $x \rightarrow E$ is invertible. For every feasible σ , we have $\langle \sigma \rangle_{x \rightarrow E} = \{t_{x \rightarrow E}^{-1}(\sigma)\}$.

Proof: Invertability tells us that $\langle \sigma \rangle_{x \rightarrow E}$ has only one element. The function $t_{x \rightarrow E}^{-1}$, given the feasible σ , produces an element of $\langle \sigma \rangle_{x \rightarrow E}$, as $(t_{x \rightarrow E}^{-1}(\sigma))[x \rightarrow E] = \sigma$. ■

Definition 55. We define an alternate means of assignment, $\delta \langle x \rightarrow E \rangle$. Let $V = \text{fv}(\delta) - \{x\}$.

- If $x \rightarrow E$ is invertible, then

$$\delta \langle x \rightarrow E \rangle = \lambda \sigma. \text{ if } \sigma \text{ is feasible} \\ \text{then } \delta(t_{x \rightarrow E}^{-1}(\sigma)) \\ \text{else } 0$$

- If $x \rightarrow E$ is not invertible, then

$$\delta \langle x \rightarrow E \rangle = \lambda \sigma. \text{ if } \sigma \text{ is feasible} \\ \text{then } \delta \upharpoonright V(\sigma \upharpoonright V) \\ \text{else } 0$$

Lemma 56. For any δ , $\delta[x \rightarrow E] = \delta \langle x \rightarrow E \rangle$.

Proof: Let $\delta' = \delta[x \rightarrow E]$ and $\delta'' = \delta \langle x \rightarrow E \rangle$.

$$\begin{aligned} \delta'(\sigma) &= \sum_{\tau \upharpoonright \tau[x \rightarrow E] = \sigma} \delta(\tau) \\ &= \sum_{\tau \in \langle \sigma \rangle_{x \rightarrow E}^V} \delta(\tau) \end{aligned}$$

Case 1: $x \rightarrow E$ is invertible

If σ is feasible, $\langle \sigma \rangle_{x \rightarrow E}$ has only one element, $\sigma^{-1} = t_{x \rightarrow E}^{-1}(\sigma)$, by Remark 54. So $\delta'(\sigma) = \delta(\sigma^{-1}) = \delta''(\sigma)$. Note that when σ^{-1} is not in $\text{support}(\delta)$ then $\delta'(\sigma) = 0 = \delta''(\sigma)$.

If σ is not feasible then $\langle \sigma \rangle_{x \rightarrow E} = \emptyset$ so $\delta'(\sigma) = 0 = \delta''(\sigma)$.

Case 2: $x \rightarrow E$ is non-invertible

If σ is feasible, then by Lemma 53 we have $\langle \sigma \rangle_{x \rightarrow E}^V = [\sigma \upharpoonright V]_{\delta}^V$.

$$\begin{aligned}
\delta'(\sigma) &= \sum_{\tau \in \langle \sigma \upharpoonright V \rangle_{\delta}^{x \rightarrow E}} \delta(\tau) \\
&= \sum_{\tau \in [\sigma \upharpoonright V]_{\delta}^V} \delta(\tau) \\
&= (\delta \upharpoonright V)(\sigma \upharpoonright V) \\
&= \delta''(\sigma)
\end{aligned}$$

If σ is not feasible then $\langle \sigma \upharpoonright V \rangle_{\delta}^{x \rightarrow E} = \emptyset$ so $\delta'(\sigma) = 0 = \delta''(\sigma)$. ■

Lemma 57. *Assume $x \rightarrow E$ is invertible, then $\text{support}(\delta) = \{t_{x \rightarrow E}^{-1}(\sigma) \mid \sigma \in \text{support}(\delta \langle x \rightarrow E \rangle)\}$.*

Proof: Let $\delta_2 = \delta \langle x \rightarrow E \rangle$. Let $\tau \in \text{support}(\delta)$. So $\sigma \stackrel{\text{def}}{=} \tau[x \rightarrow E] \in \text{support}(\delta_2)$ and $t_{x \rightarrow E}^{-1}(\sigma) = \tau$. So $\tau \in \{t_{x \rightarrow E}^{-1}(\sigma) \mid \sigma \in \text{support}(\delta_2)\}$.

Let $\tau \in \{t_{x \rightarrow E}^{-1}(\sigma) \mid \sigma \in \text{support}(\delta_2)\}$. So $\tau = t_{x \rightarrow E}^{-1}(\sigma)$ for some $\sigma \in \text{support}(\delta_2)$. So there exists $\tau' \in \text{support}(\delta)$ such that $\tau'[x \rightarrow E] = \sigma$. But $t_{x \rightarrow E}^{-1}(\sigma) = t_{x \rightarrow E}^{-1}(t_{x \rightarrow E}(\tau')) = \tau'$ so $\tau' = \tau$ as $\tau = t_{x \rightarrow E}^{-1}(\sigma)$. So $\tau \in \text{support}(\delta)$. ■

Lemma 9 (Soundness of Assignment). *If $\delta \in \gamma_{\mathbb{P}}(P)$ then $\delta[x \rightarrow E] \in \gamma_{\mathbb{P}}(P[x \rightarrow E])$.*

Proof: Let $V = \text{fv}(\delta) - \{x\}$. By assumption, we have the following.

$$\text{support}(\delta) \subseteq \gamma_{\mathbb{C}}(C) \quad (24)$$

$$s^{\min} \leq |\text{support}(\delta)| \leq s^{\max} \quad (25)$$

$$m^{\min} \leq \|\delta\| \leq m^{\max} \quad (26)$$

$$\forall \sigma \in \text{support}(\delta) . p^{\min} \leq \delta(\sigma) \leq p^{\max} \quad (27)$$

Let $P_2 = P[x \rightarrow E]$ and $\delta_2 = \delta[x \rightarrow E] = \delta \langle x \rightarrow E \rangle$. Lemma 56 lets us use $\delta[x \rightarrow E]$ or $\delta \langle x \rightarrow E \rangle$ interchangeably.

We consider two cases. **Case 1: $x \rightarrow E$ is invertible**

In this case, P_2 is defined with $C_2 = C[x \rightarrow E]$ and all other parameters as in P . Thus we need to show the following.

$$\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2) \quad (28)$$

$$s_1^{\min} = s_2^{\min} \leq |\text{support}(\delta_2)| \leq s_2^{\max} = s_1^{\max} \quad (29)$$

$$m_1^{\min} = m_2^{\min} \leq \|\delta_2\| \leq m_2^{\max} = m_1^{\max} \quad (30)$$

$$\forall \sigma \in \text{support}(\delta_2) . p_1^{\min} = p_2^{\min} \leq \delta_2(\sigma) \leq p_2^{\max} = p_1^{\max} \quad (31)$$

Claim (28) Support. By definition, $\gamma_{\mathbb{C}}(C_2) = \{\sigma[x \rightarrow E] \mid \sigma \in \gamma_{\mathbb{C}}(C)\}$. Let $\tau \in \text{support}(\delta_2)$, so we have $\sigma \in \text{support}(\delta) \subseteq \gamma_{\mathbb{C}}(C)$ with $\sigma[x \rightarrow E] = \tau$. So $\tau \in \gamma_{\mathbb{C}}(C_2)$. So $\tau \in \gamma_{\mathbb{C}}(C_2)$ and thus $\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2)$. □

Claim (29) Support points. By Lemma 57 we have $\text{support}(\delta) = \{t_{x \rightarrow E}^{-1}(\sigma) \mid \sigma \in \text{support}(\delta_2)\}$. Inverse functions are necessarily injective over their domain, and since $\text{support}(\delta_2)$ are all feasible (thus in the domain of the inverse), we have $|\{t_{x \rightarrow E}^{-1}(\sigma) \mid \sigma \in \text{support}(\delta_2)\}| = |\text{support}(\delta_2)|$. So $|\text{support}(\delta)| = |\text{support}(\delta_2)|$. This, together with (25), completes the claim. □

Claim (30) Mass. Note again that $\text{support}(\delta_2) \subseteq \text{State}[x \rightarrow E]$. That is, all possible states are feasible. So we can write:

$$\begin{aligned}
\|\delta_2\| &= \sum_{\sigma \in \text{support}(\delta_2)} \delta_2(\sigma) \\
&= \sum_{\sigma \in \text{support}(\delta_2)} \delta(t_{x \rightarrow E}^{-1}(\sigma)) \quad [\text{by defn. of } \delta_2] \\
&= \sum_{\tau \in \text{support}(\delta)} \delta(\tau) \quad [\text{by Lemma 57}] \\
&= \|\delta\|
\end{aligned}$$

The above, together with (26), completes this claim. □

Claim (31) Probability. Since $\text{support}(\delta_2)$ are feasible, we have, for every $\sigma \in \text{support}(\delta_2)$, $\delta_2(\sigma) = \delta(t_{x \rightarrow E}^{-1}(\sigma))$. But also, $t_{x \rightarrow E}^{-1}(\sigma) \in \text{support}(\delta)$. Taking this, and (27), completes this claim, and soundness in the invertible case. □

Case 2: $x \rightarrow E$ is non-invertible In this case, P_2 is defined via the forget operation. If $P_1 = f_x(P)$ and $C_1 = (B_1, V_1)$, then $P_2 = P[x \rightarrow E]$ has $C_2 = (B_1 \cup \{x = E\}, V_1 \cup \{x\})$, and all other parameters as in P_1 .

We need to show the following four claims.

$$\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2) \quad (32)$$

$$s_1^{\min} = s_2^{\min} \leq |\text{support}(\delta_2)| \leq s_2^{\max} = s_1^{\max} \quad (33)$$

$$m_1^{\min} = m_2^{\min} \leq \|\delta_2\| \leq m_2^{\max} = m_1^{\max} \quad (34)$$

$$\forall \sigma \in \text{support}(\delta_2) . p_1^{\min} = p_2^{\min} \leq \delta_2(\sigma) \leq p_2^{\max} = p_1^{\max} \quad (35)$$

Recall the definition of δ_2 :

$$\begin{aligned}
\delta \langle x \rightarrow E \rangle &= \lambda \sigma . \text{if } \sigma \text{ is feasible} \\
&\quad \text{then } \delta \upharpoonright V(\sigma \upharpoonright V) \\
&\quad \text{else } 0
\end{aligned}$$

Claim (32) Support. Let $\sigma \in \text{support}(\delta_2)$. So $\sigma \upharpoonright V \in \text{support}(\delta \upharpoonright V)$. so there exists $\tau \in \text{support}(\delta) \subseteq \gamma_{\mathbb{C}}(C)$ with $\tau \upharpoonright V = \sigma \upharpoonright V$. So $\tau \upharpoonright V \in \gamma_{\mathbb{C}}(f_x(C)) = \gamma_{\mathbb{C}}(C_1)$. So $\tau \in \gamma_{\mathbb{C}}((B_1, V_1 \cup \{x\}))$ as the add dimension operation leaves x unconstrained. The non-constraint of x also tells us that $\sigma \in \gamma_{\mathbb{C}}((B_1, V_1 \cup \{x\}))$ as we have $\sigma \upharpoonright V = \tau \upharpoonright V$.

Since $\sigma \in \text{support}(\delta_2)$, σ is feasible so σ satisfies the $x = E$ constraint as $\sigma = \tau[x \rightarrow E]$ for some τ . Thus, overall, we have $\sigma \in \gamma_{\mathbb{C}}((B_1 \cup \{x = E\}, V_1 \cup \{x\})) = \gamma_{\mathbb{C}}(C_2)$. □

Claim (33) Support points. Let $\delta_1 = \delta \upharpoonright V = f_x(\delta)$. By soundness of forget (Lemma 7), we have the following. □

$$s_1^{\min} = s_2^{\min} \leq |\text{support}(\delta_1)| \leq s_2^{\max} = s_1^{\max}$$

All we need to show, then, is the following.

$$|\text{support}(\delta_1)| = |\text{support}(\delta_2)| \quad (36)$$

Let us show this by establishing a bijection f between the two sets. Let us define $f : \text{support}(\delta_1) \rightarrow \text{support}(\delta_2)$ via $f : \sigma_V \mapsto \sigma_V \cup \{x = \llbracket E \rrbracket \sigma_V\}$.

To show f is injective, let σ_V, σ'_V be such that $f(\sigma_V) = f(\sigma'_V)$. Since f does not change any part of the state other than adding x , it must be that $\sigma_V = \sigma'_V$.

To show that f is surjective, consider $\sigma \in \text{support}(\delta_2)$. So σ is feasible, so $\sigma[x \rightarrow E] = \sigma$ by Lemma 52. Also $\sigma \upharpoonright V \in \text{support}(\sigma_V)$, considering the definition of σ_2 . Since E doesn't depend on x , we can write $\llbracket E \rrbracket \sigma = \llbracket E \rrbracket \sigma \upharpoonright V$, therefore $f(\sigma \upharpoonright V) = \sigma \upharpoonright V \cup \{x = \llbracket E \rrbracket \sigma \upharpoonright V\} = \sigma[x \rightarrow E] = \sigma$.

Since f is injective and surjective, it is a bijection and thus $|\text{support}(\delta_1)| = |\text{support}(\delta_2)|$. □

Claim (34) Mass. Let $\delta_1 = \delta \upharpoonright V = f_x(\delta)$. Let us show the following claim.

$$LHS = \text{support}(\delta_1) = \{\sigma \upharpoonright V \mid \sigma \in \text{support}(\delta_2)\} = RHS \quad (37)$$

Let $\sigma_V \in \text{support}(\delta_1)$. So there exists $\sigma \in \text{support}(\delta)$ with $\sigma \upharpoonright V = \sigma_V$. So $\sigma[x \rightarrow E] \in \text{support}(\delta_2)$. But the assignment doesn't change anything but x , so it must be that $(\sigma[x \rightarrow E]) \upharpoonright V = \sigma \upharpoonright V$, therefore $\sigma_V = \sigma \upharpoonright V \in \{\tau \upharpoonright V \mid \tau \in \text{support}(\delta_2)\}$. Thus $LHS \subseteq RHS$.

On the other side, let $\sigma \in \text{support}(\delta_2)$, so $\sigma = \tau[x \rightarrow E]$ for some $\tau \in \text{support}(\delta)$, by the original definition of distribution assignment. So $\tau \upharpoonright V \in \text{support}(\delta_1)$. But $(\tau[x \rightarrow E]) \upharpoonright V = \tau \upharpoonright V$ as the assignment doesn't change anything but x . So $\sigma \upharpoonright V = (\tau[x \rightarrow E]) \upharpoonright V = \tau \upharpoonright V \in \text{support}(\delta_1)$, concluding that $RHS \subseteq LHS$, and thus $LHS = RHS$.

Note that this, together with (36), show that not only are the sets equal, but also no two elements of $\text{support}(\delta_2)$ can map, via projection to V , to the same element of $\text{support}(\delta_1)$.

By soundness of forget (Lemma 7), we have the following.

$$m_1^{\min} = m_2^{\min} \leq \|\delta_1\| \leq m_2^{\max} = m_1^{\max}$$

Again, we proceed to show that $\|\delta_1\| = \|\delta_2\|$.

$$\begin{aligned} \|\delta_1\| &= \sum_{\sigma_V \in \text{support}(\delta_1)} \delta_1(\sigma_V) \\ &= \sum_{\sigma \in \text{support}(\delta_2)} \delta_1(\sigma \upharpoonright V) \quad [\text{by (36) and (37)}] \\ &= \sum_{\sigma \in \text{support}(\delta_2)} \delta_2(\sigma) \quad [\text{by defn. of } \delta_2] \\ &= \|\delta_2\| \end{aligned}$$

Claim (35) Probability. Let $\sigma \in \text{support}(\delta_2)$. So σ is feasible, so $\delta_2(\sigma) = (\delta \upharpoonright V)(\sigma \upharpoonright V) > 0$. Therefore $\sigma \upharpoonright V \in \text{support}(\delta \upharpoonright V)$. Thus, by soundness of forget (Lemma 7), we have $p_2^{\min} = p_1^{\min} \leq (\delta \upharpoonright V)(\sigma \upharpoonright V) \leq p_1^{\max} = p_2^{\max}$, concluding the claim and the lemma. □

C. Plus

Definition 58. Let $\text{overlap}(\delta_1, \delta_2) = \text{support}(\delta_1) \cap \text{support}(\delta_2)$.

Lemma 59. If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$ then $P_1 \odot P_2 \leq |\text{overlap}(\delta_1, \delta_2)| \leq P_1 \odot P_2$.

Proof: We first note that for any sets A, B , it is the case that $|A \cup B| = |A| + |B| - |A \cap B|$ (often called the ‘‘inclusion-exclusion principle’’). Rearranging the equation we also have $|A \cap B| = |A| + |B| - |A \cup B|$.

We will make use of this formula with $A = \text{support}(\delta_1)$, $B = \text{support}(\delta_2)$.

Lower Bound: We first show the lower bound. Expanding the definitions of $P_1 \odot P_2$ and $\text{overlap}(\delta_1, \delta_2)$, this reduces to showing the following.

$$\begin{aligned} \max((s_1^{\min} - n_1) + (s_2^{\min} - n_2) - n_3, 0) \\ \leq |\text{support}(\delta_1) \cap \text{support}(\delta_2)| \end{aligned}$$

Clearly we have $0 \leq |\text{support}(\delta_1) \cap \text{support}(\delta_2)|$, so it remains to show that the following holds.

$$\begin{aligned} (s_1^{\min} - n_1) + (s_2^{\min} - n_2) - n_3 \\ \leq |\text{support}(\delta_1) \cap \text{support}(\delta_2)| \end{aligned}$$

Expanding the definitions of n_1, n_2 from Definition 10, we obtain

$$\begin{aligned} (s_1^{\min} - (\#(C_1) - n_3)) + (s_2^{\min} - (\#(C_2) - n_3)) - n_3 \\ \leq |\text{support}(\delta_1) \cap \text{support}(\delta_2)| \end{aligned}$$

and rearranging yields the following.

$$\begin{aligned} s_1^{\min} + s_2^{\min} - (\#(C_1) + \#(C_2) - n_3) \\ \leq |\text{support}(\delta_1) \cap \text{support}(\delta_2)| \end{aligned}$$

This follows from the rearranged inclusion-exclusion principle provided we can show $s_1^{\min} \leq |\text{support}(\delta_1)|$, $s_2^{\min} \leq |\text{support}(\delta_2)|$, and $\#(C_1) + \#(C_2) - n_3 \geq |\text{support}(\delta_1) \cup \text{support}(\delta_2)|$. The first two follow directly from our assumptions that $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$. For the third condition, we reason as follows.

We have from our assumptions that $\gamma_{\mathbb{C}}(C_1) \supseteq \text{support}(\delta_1)$ and $\gamma_{\mathbb{C}}(C_2) \supseteq \text{support}(\delta_2)$. Thus, we have

$$\gamma_{\mathbb{C}}(C_1) \cup \gamma_{\mathbb{C}}(C_2) \supseteq \text{support}(\delta_1) \cup \text{support}(\delta_2)$$

and finally

$$|\gamma_{\mathbb{C}}(C_1) \cup \gamma_{\mathbb{C}}(C_2)| \geq |\text{support}(\delta_1) \cup \text{support}(\delta_2)|$$

Utilizing the inclusion-exclusion principle, we have

$$\begin{aligned} |\gamma_{\mathbb{C}}(C_1)| + |\gamma_{\mathbb{C}}(C_2)| - |\gamma_{\mathbb{C}}(C_1) \cap \gamma_{\mathbb{C}}(C_2)| \\ \geq |\text{support}(\delta_1) \cup \text{support}(\delta_2)| \end{aligned}$$

Since we have $|\gamma_{\mathbb{C}}(C)| = \#(C)$, we can rewrite this to the following.

$$\begin{aligned} \#(C_1) + \#(C_2) - |\gamma_{\mathbb{C}}(C_1) \cap \gamma_{\mathbb{C}}(C_2)| \\ \geq |\text{support}(\delta_1) \cup \text{support}(\delta_2)| \end{aligned}$$

It remains to show that $|\gamma_{\mathbb{C}}(C_1) \cap \gamma_{\mathbb{C}}(C_2)| = n_3$. We have that $\gamma_{\mathbb{C}}(C_1 \sqcap_{\mathbb{C}} C_2) = \gamma_{\mathbb{C}}(C_1) \cap \gamma_{\mathbb{C}}(C_2)$ (that is, $\sqcap_{\mathbb{C}}$ is precise). This allows us to complete the final step, concluding that n_3 , which is defined as $\#(C_1 \sqcap_{\mathbb{C}} C_2)$ is equal to $|\gamma_{\mathbb{C}}(C_1) \cap \gamma_{\mathbb{C}}(C_2)|$.

Upper Bound: We next show that the upper bound holds. Our goal is to show the following.

$$P_1 \odot P_2 \geq |\text{overlap}(\delta_1, \delta_2)|$$

Expanding our definitions yields the following formula.

$$\min(s_1^{\max}, s_2^{\max}, n_3) \geq |\text{support}(\delta_1) \cap \text{support}(\delta_2)|$$

We first note that the following holds.

$$|\text{support}(\delta_1) \cap \text{support}(\delta_2)| \leq |\text{support}(\delta_1)| \leq s_1^{\max}$$

Thus s_1^{\max} is a sound upper bound. Similarly, we have

$$|\text{support}(\delta_1) \cap \text{support}(\delta_2)| \leq |\text{support}(\delta_2)| \leq s_2^{\max}$$

which shows that s_2^{\max} is a sound upper bound. Finally, we note that our assumptions give us $\text{support}(\delta_1) \subseteq \gamma_{\mathbb{C}}(C_1)$ and $\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2)$. Thus we have the following.

$$\text{support}(\delta_1) \cap \text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_1) \cap \gamma_{\mathbb{C}}(C_2)$$

We showed previously that $n_3 = |\gamma_{\mathbb{C}}(C_1) \cap \gamma_{\mathbb{C}}(C_2)|$. Thus we have

$$|\text{support}(\delta_1) \cap \text{support}(\delta_2)| \leq n_3$$

which shows that n_3 is a sound upper bound.

Since all of s_1^{\max} , s_2^{\max} , and n_3 are sound upper bounds, their minimum is also a sound upper bound. ■

Lemma 60.

$$\begin{aligned} |\text{support}(\delta_1 + \delta_2)| = \\ |\text{support}(\delta_1)| + |\text{support}(\delta_2)| - |\text{overlap}(\delta_1, \delta_2)| \end{aligned}$$

Proof: First we note that $\text{support}(\delta_1 + \delta_2) = \{\sigma \mid \delta_1(\sigma) + \delta_2(\sigma) > 0\}$. Since the range of δ_1 and δ_2 is $[0, 1]$, we

have that $\delta_1(\sigma) + \delta_2(\sigma) > 0$ if and only if either $\delta_1(\sigma) > 0$ or $\delta_2(\sigma) > 0$. Thus, we have $\sigma \in \text{support}(\delta_1 + \delta_2)$ if and only if $\sigma \in \text{support}(\delta_1)$ or $\sigma \in \text{support}(\delta_2)$, which implies $\text{support}(\delta_1 + \delta_2) = \text{support}(\delta_1) \cup \text{support}(\delta_2)$.

Next, we note that for any sets A, B we have $|A \cup B| = |A| + |B| - |A \cap B|$. Utilizing this statement with $A = \text{support}(\delta_1)$ and $B = \text{support}(\delta_2)$ completes the proof. ■

Lemma 12 (Soundness of Plus). *If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$ then $\delta_1 + \delta_2 \in \gamma_{\mathbb{P}}(P_1 + P_2)$.*

Proof: Suppose $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$. Then we have the following.

$$\text{support}(\delta_1) \subseteq \gamma_{\mathbb{C}}(C_1) \quad (38)$$

$$s_1^{\min} \leq |\text{support}(\delta_1)| \leq s_1^{\max} \quad (39)$$

$$m_1^{\min} \leq \|\delta_1\| \leq m_1^{\max} \quad (40)$$

$$\forall \sigma \in \text{support}(\delta_1). p_1^{\min} \leq \delta_1(\sigma) \leq p_1^{\max} \quad (41)$$

and

$$\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2) \quad (42)$$

$$s_2^{\min} \leq |\text{support}(\delta_2)| \leq s_2^{\max} \quad (43)$$

$$m_2^{\min} \leq \|\delta_2\| \leq m_2^{\max} \quad (44)$$

$$\forall \sigma \in \text{support}(\delta_2). p_2^{\min} \leq \delta_2(\sigma) \leq p_2^{\max} \quad (45)$$

The definition of abstract plus has special cases when either of the arguments are zero, that is, if $\text{iszero}(P_1)$ or $\text{iszero}(P_2)$. Without the loss of generality, let us assume $\text{iszero}(P_2)$ and thus by definition $P_1 + P_2 = P_1$. Since $\gamma_{\mathbb{P}}(P_2) = \{0_{\text{Dist}}\}$, where 0_{Dist} is the distribution assigning probability of 0 to every state. Therefore $\delta_2 = 0_{\text{Dist}}$ and thus $\delta_1 + \delta_2 = \delta_1$. But we already have $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ by assumption, hence we are done in this case.

In the case when not $\text{iszero}(P_1)$ and not $\text{iszero}(P_2)$ we must show the following.

$$\text{support}(\delta_1 + \delta_2) \subseteq \gamma_{\mathbb{C}}(C_1 \sqcup_{\mathbb{C}} C_2) \quad (46)$$

$$\max\{s_1^{\min} + s_2^{\min} - P_1 \odot P_2, 0\} \leq |\text{support}(\delta_1 + \delta_2)| \quad (47)$$

$$|\text{support}(\delta_1 + \delta_2)| \leq \min\{s_1^{\max} + s_2^{\max} - P_1 \odot P_2, \#(C_3)\} \quad (48)$$

$$m_1^{\min} + m_2^{\min} \leq \|\delta_1 + \delta_2\| \leq m_1^{\max} + m_2^{\max} \quad (49)$$

We also must show the conditions on p^{\min} and p^{\max} for the sum.

Condition (46) follows from (38) and (42) and the fact that $\sqcup_{\mathbb{C}}$ over-approximates union. The key step is noting that $\text{support}(\delta_1 + \delta_2) = \text{support}(\delta_1) \cup \text{support}(\delta_2)$. To show this we consider some $\sigma \in \text{support}(\delta_1 + \delta_2)$. We have that $(\delta_1 + \delta_2)(\sigma) > 0$ which, expanding the definition of $+$, yields $\delta_1(\sigma) + \delta_2(\sigma) > 0$. Since the range of δ_1 and δ_2 is $[0, 1]$, this implies that either $\delta_1(\sigma) > 0$ or $\delta_2(\sigma) > 0$ and thus $\sigma \in \text{support}(\delta_1)$ or $\sigma \in \text{support}(\delta_2)$.

Conditions (47) and (48) follow from (39) and (43) and Lemmas 59 and 60. We have $s_1^{\min} \leq |\text{support}(\delta_1)|$ from (39) and $s_2^{\min} \leq |\text{support}(\delta_2)|$ from (43). Monotonicity of addition then gives us

$$s_1^{\min} + s_2^{\min} \leq |\text{support}(\delta_1)| + |\text{support}(\delta_2)|$$

From Lemma 59 we have $|\text{overlap}(\delta_1, \delta_2)| \leq P_1 \odot P_2$ and thus

$$-P_1 \odot P_2 \leq -|\text{overlap}(\delta_1, \delta_2)|$$

Combining with the above yields

$$s_1^{\min} + s_2^{\min} - P_1 \odot P_2 \leq |\text{support}(\delta_1)| + |\text{support}(\delta_2)| - |\text{overlap}(\delta_1, \delta_2)|$$

We can then rewrite the right-hand side according to Lemma 60 to obtain

$$s_1^{\min} + s_2^{\min} - P_1 \odot P_2 \leq |\text{support}(\delta_1 + \delta_2)|$$

which is condition (47).

Condition (48) follows the same reasoning. We have $|\text{support}(\delta_1)| + |\text{support}(\delta_2)| \leq s_1^{\max} + s_2^{\max}$ by (39) and (43). We then apply Lemma 59 and 60 to obtain condition (48).

For Condition (49), note that

$$\|\delta_1 + \delta_2\| = \sum_{\sigma} (\delta_1(\sigma) + \delta_2(\sigma)) = \sum_{\sigma} \delta_1(\sigma) + \sum_{\sigma} \delta_2(\sigma)$$

This is then equivalent to $\|\delta_1\| + \|\delta_2\|$. We have shown that $\|\delta_1 + \delta_2\| = \|\delta_1\| + \|\delta_2\|$. Condition (49) then follows from monotonicity of addition applied to (40) and (44)

We now consider the p^{\min} and p^{\max} conditions. Let $P_3 = P_1 + P_2$ and $\delta_3 = \delta_1 + \delta_2$. We must show.

$$\forall \sigma \in \text{support}(\delta_3) . p_3^{\min} \leq \delta_3(\sigma) \leq p_3^{\max}$$

The values p_3^{\min} and p_3^{\max} are defined by cases and we consider these cases separately. In one case, we have that p^{\min} of the sum is $\min(p_1^{\min}, p_2^{\min})$. This is always a sound choice. To see why, suppose $\sigma \in \text{support}(\delta_1 + \delta_2)$. Then $\sigma \in \text{support}(\delta_1)$ or $\sigma \in \text{support}(\delta_2)$. If $\sigma \in \text{support}(\delta_1)$, then $(\delta_1 + \delta_2)(\sigma) = \delta_1(\sigma) + \delta_2(\sigma)$ is at least p_1^{\min} . Similarly, if $\sigma \in \text{support}(\delta_2)$ then $(\delta_1 + \delta_2)(\sigma) \geq \delta_2(\sigma)$.

Similarly, the value $p_1^{\max} + p_2^{\max}$ is always a sound choice for p_3^{\max} . Consider $\sigma \in \text{support}(\delta_3)$. Then $\sigma \in \text{support}(\delta_1)$ or $\sigma \in \text{support}(\delta_2)$. If $\sigma \in \text{support}(\delta_1)$ and $\sigma \notin \text{support}(\delta_2)$, then we have

$$\delta_3(\sigma) = \delta_1(\sigma) + \delta_2(\sigma) = \delta_1(\sigma)$$

By (41) we then have $\delta_3(\sigma) \leq p_1^{\max}$ and thus $\delta_3(\sigma) \leq p_1^{\max} + p_2^{\max}$ as desired.

Similarly, if $\sigma \notin \text{support}(\delta_1)$ and $\sigma \in \text{support}(\delta_2)$ then by (45) we have

$$\delta_3(\sigma) = \delta_2(\sigma) \leq p_2^{\max} \leq p_1^{\max} + p_2^{\max}$$

Finally, if $\sigma \in \text{support}(\delta_1)$ and $\sigma \in \text{support}(\delta_2)$ then by (41) we have $\delta_1(\sigma) \leq p_1^{\max}$. By (45) we have $\delta_2(\sigma) \leq p_2^{\max}$. Combining these we have $\delta_1(\sigma) + \delta_2(\sigma) \leq p_1^{\max} + p_2^{\max}$ which is equivalent to $\delta_3(\sigma) \leq p_3^{\max}$ as desired.

Next we consider the $P_1 \odot P_2 = \#(C_3)$ case for p_3^{\min} . We must show that $p_1^{\min} + p_2^{\min}$ is a sound lower bound on $\delta_3(\sigma)$ for $\sigma \in \text{support}(\delta_3)$. We have by Lemma 59 that $P_1 \odot P_2 \leq |\text{overlap}(\delta_1, \delta_2)|$. Since $P_1 \odot P_2 = \#(C_3)$ and $\#(C_3) \geq |\text{overlap}(\delta_1, \delta_2)|$, we have that $\#(C_3) = |\text{overlap}(\delta_1, \delta_2)|$. Expanding the definition of $\text{overlap}(\delta_1, \delta_2)$ yields

$$|\text{support}(\delta_1) \cap \text{support}(\delta_2)| = \#(C_3) \quad (50)$$

We have from (46) that $\text{support}(\delta_1 + \delta_2) \subseteq \gamma_{\mathbb{C}}(C_3)$ and from the proof of (46) we have that $\text{support}(\delta_1 + \delta_2) = \text{support}(\delta_1) \cup \text{support}(\delta_2)$. Combining these yields

$$|\text{support}(\delta_1) \cup \text{support}(\delta_2)| \leq \#(C_3)$$

Combining this with (50) yields

$$|\text{support}(\delta_1) \cup \text{support}(\delta_2)| \leq |\text{support}(\delta_1) \cap \text{support}(\delta_2)|$$

For any sets A, B , we have that $|A \cup B| \geq |A \cap B|$ and thus the above inequality implies the following.

$$|\text{support}(\delta_1) \cup \text{support}(\delta_2)| = |\text{support}(\delta_1) \cap \text{support}(\delta_2)|$$

The fact that the size of the intersection and the size of the union of $\text{support}(\delta_1)$ and $\text{support}(\delta_2)$ is identical implies that $\text{support}(\delta_1) = \text{support}(\delta_2)$. This implies that for all σ , we have $\sigma \in \text{support}(\delta_1)$ if and only if $\sigma \in \text{support}(\delta_2)$.

Now consider $\sigma \in \text{support}(\delta_3)$. We have $\sigma \in \text{support}(\delta_1)$ or $\sigma \in \text{support}(\delta_2)$, as before, but now we can strengthen this to $\sigma \in \text{support}(\delta_1)$ and $\sigma \in \text{support}(\delta_2)$. By (41) we have $p_1^{\min} \leq \delta_1(\sigma)$ and by (45) we have $p_2^{\min} \leq \delta_2(\sigma)$. Thus we have

$$p_1^{\min} + p_2^{\min} \leq \delta_1(\sigma) + \delta_2(\sigma)$$

which was our goal.

Finally we consider the $P_1 \odot P_2 = 0$ case for p_3^{\max} (the ‘‘otherwise’’ case in Definition 11). Consider a $\sigma \in \text{support}(\delta_3)$. We must show that $\delta_3(\sigma) \leq \max(p_1^{\max}, p_2^{\max})$. We have that either $\sigma \in \text{support}(\delta_1)$ or $\sigma \in \text{support}(\delta_2)$. We cannot have both since $P_1 \odot P_2 = 0$ which, by Lemma 59 implies that $|\text{overlap}(\delta_1, \delta_2)| = 0$. If $\sigma \in \text{support}(\delta_1)$ then by (41) we have $\delta_1(\sigma) \leq p_1^{\max}$. We have $\sigma \notin \text{support}(\delta_2)$ and thus $\delta_2(\sigma) = 0$. Thus we reason that

$$\delta_1(\sigma) + \delta_2(\sigma) = \delta_1(\sigma) \leq p_1^{\max} \leq \max(p_1^{\max}, p_2^{\max})$$

Similarly, if $\sigma \in \text{support}(\delta_2)$ then we apply (45) to obtain

$$\delta_1(\sigma) + \delta_2(\sigma) = \delta_2(\sigma) \leq p_2^{\max} \leq \max(p_1^{\max}, p_2^{\max})$$

■

D. Product

Lemma 13 (Soundness of Product). *If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ and $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$ then $\delta_1 \times \delta_2 \in \gamma_{\mathbb{P}}(P_1 \times P_2)$.*

Proof: By assumption, we have the following for $i = 1, 2$.

$$\text{support}(\delta_i) \subseteq \gamma_{\mathbb{C}}(C_i) \quad (51)$$

$$s_i^{\min} \leq |\text{support}(\delta_i)| \leq s_i^{\max} \quad (52)$$

$$m_i^{\min} \leq \|\delta_i\| \leq m_i^{\max} \quad (53)$$

$$\forall \sigma \in \text{support}(\delta_i) \cdot p_i^{\min} \leq \delta(\sigma_i) \leq p_i^{\max} \quad (54)$$

Let $\delta_3 = \delta_1 \times \delta_2$ and $P_3 = P_1 \times P_2$. Recall the definition of P_3 .

$$\begin{array}{l} C_3 = C_1 \times C_2 \\ \left. \begin{array}{l} p_3^{\min} = p_1^{\min} \cdot p_2^{\min} \\ s_3^{\min} = s_1^{\min} \cdot s_2^{\min} \\ m_3^{\min} = m_1^{\min} \cdot m_2^{\min} \end{array} \right| \begin{array}{l} p_3^{\max} = p_1^{\max} \cdot p_2^{\max} \\ s_3^{\max} = s_1^{\max} \cdot s_2^{\max} \\ m_3^{\max} = m_1^{\max} \cdot m_2^{\max} \end{array} \end{array}$$

We must show the following four claims.

$$\text{support}(\delta_3) \subseteq \gamma_{\mathbb{C}}(C_3) \quad (55)$$

$$s_3^{\min} \leq |\text{support}(\delta_3)| \leq s_3^{\max} \quad (56)$$

$$m_3^{\min} \leq \|\delta_3\| \leq m_3^{\max} \quad (57)$$

$$\forall \sigma \in \text{support}(\delta_3) \cdot p_3^{\min} \leq \delta_3(\sigma) \leq p_3^{\max} \quad (58)$$

Also, recall the definition of concrete product.

$$\delta_1 \times \delta_2 = \lambda(\sigma_1, \sigma_2) \cdot \delta_1(\sigma_1) \cdot \delta_2(\sigma_2)$$

Let $V_1 = \text{fv}(\delta_1)$ and $V_2 = \text{fv}(\delta_2)$.

Claim (55) – Support. Let $\sigma = (\sigma_1, \sigma_2) \in \text{support}(\delta_3)$. Thus it must be that $\delta_1(\sigma_1) > 0$ and $\delta_2(\sigma_2) > 0$, thus, by (51), $\sigma_1 \in \text{support}(\delta_1) \subseteq \gamma_{\mathbb{C}}(C_1)$ and $\sigma_2 \in \text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2)$, therefore $\sigma \in \gamma_{\mathbb{C}}(\delta_3)$. \square

Claim (56) – Support points. Using (52) we get the following.

$$s_1^{\min} \cdot s_2^{\min} \leq |\text{support}(\delta_1)| \cdot |\text{support}(\delta_2)| \leq s_1^{\max} \cdot s_2^{\max}$$

Likewise, the size of $\text{support}(\delta_3)$ can be equated as follows.

$$\begin{aligned} |\text{support}(\delta_3)| &= \left| \left\{ (\sigma_1, \sigma_2) \mid \begin{array}{l} \sigma_1 \in \text{support}(\delta_1), \\ \sigma_2 \in \text{support}(\delta_2) \end{array} \right\} \right| \\ &= |\text{support}(\delta_1)| \cdot |\text{support}(\delta_2)| \end{aligned}$$

This completes the claim as $s_3^{\min} = s_1^{\min} \cdot s_2^{\min}$ and $s_3^{\max} = s_1^{\max} \cdot s_2^{\max}$. \square

Claim (57) – Mass.

$$\begin{aligned} \|\delta_3\| &= \sum_{\sigma \in \text{support}(\delta_3)} \delta(\sigma) \\ &= \sum_{(\sigma_1, \sigma_2) \in \text{support}(\delta_3)} \delta_1(\sigma_1) \cdot \delta_2(\sigma_2) \\ &= \sum_{\sigma_1 \in \text{support}(\delta_1)} \left(\sum_{\sigma_2 \in \text{support}(\delta_2)} \delta_1(\sigma_1) \cdot \delta_2(\sigma_2) \right) \\ &= \sum_{\sigma_1 \in \text{support}(\delta_1)} \delta_1(\sigma_1) \sum_{\sigma_2 \in \text{support}(\delta_2)} \delta_2(\sigma_2) \\ &= \sum_{\sigma_1 \in \text{support}(\delta_1)} \delta_1(\sigma_1) \cdot \|\delta_2\| \\ &= \|\delta_1\| \cdot \|\delta_2\| \end{aligned}$$

\square

Likewise, by (53), we have the following.

$$m_1^{\min} \cdot m_2^{\min} \leq \|\delta_1\| \cdot \|\delta_2\| \leq m_1^{\max} \cdot m_2^{\max}$$

This completes the claim as $m_3^{\min} = m_1^{\min} \cdot m_2^{\min}$ and $m_3^{\max} = m_1^{\max} \cdot m_2^{\max}$.

Claim (58) – Probability. Let $\sigma = (\sigma_1, \sigma_2) \in \text{support}(\delta_3)$. Thus $\sigma_1 \in \text{support}(\delta_1)$ and $\sigma_2 \in \text{support}(\delta_2)$. Also, $\delta_3(\sigma) = \delta_1(\sigma_1) \cdot \delta_2(\sigma_2)$. By (54), we have $p_1^{\min} \leq \delta_1(\sigma_1) \leq p_1^{\max}$ and $p_2^{\min} \leq \delta_2(\sigma_2) \leq p_2^{\max}$. Therefore

$$p_3^{\min} = p_1^{\min} \cdot p_2^{\min} \leq \delta_3(\sigma) \leq p_1^{\max} \cdot p_2^{\max} = p_3^{\max}$$

This completes the claim and the proof. \square

\blacksquare

E. Conditioning

Definition 61. Given a set of states S and a boolean expression B , let $S|B$ be the subset of S that satisfy the condition B and $S|\bar{B}$ be the subset of S that do not satisfy the condition. Formally,

$$S|B \stackrel{\text{def}}{=} \{ \sigma \in S \mid \llbracket B \rrbracket \sigma = \text{true} \}$$

$$S|\bar{B} \stackrel{\text{def}}{=} \{ \sigma \in S \mid \llbracket B \rrbracket \sigma = \text{false} \}$$

Lemma 15 (Soundness of Conditioning). *If $\delta \in \gamma_{\mathbb{P}}(P)$ then $\delta|B \in \gamma_{\mathbb{P}}(P|B)$.*

Proof: Let $\delta_2 = \delta|B$. Recall the definition of the conditional distribution:

$$\delta|B = \lambda\sigma. \text{ if } \llbracket B \rrbracket \sigma \text{ then } \delta(\sigma) \text{ else } 0$$

Let $P_2 = P|B$. The construction of P_2 produces the following parameters.

$$\begin{array}{l} p_2^{\min} = p^{\min} \quad \left| \quad \begin{array}{l} s_2^{\min} = \max \{ s^{\min} - \bar{n}, 0 \} \\ p_2^{\max} = p^{\max} \quad \left| \quad \begin{array}{l} s_2^{\max} = \min \{ s^{\max}, n \} \\ m_2^{\min} = \max \{ p_2^{\min} \cdot s_2^{\min}, m^{\min} - p^{\max} \cdot \min \{ s^{\max}, \bar{n} \} \} \\ m_2^{\max} = \min \{ p_2^{\max} \cdot s_2^{\max}, m^{\max} - p^{\min} \cdot \max \{ s^{\min} - n, 0 \} \} \end{array} \right. \\ C_2 = \langle\langle B \rangle\rangle C \end{array} \right. \end{array}$$

29 The quantities n and \bar{n} are defined in such a way that n over-approximates the number of support points of δ that satisfy B , whereas \bar{n} over-approximates the number of support points of δ that do not satisfy B . Also, $\langle\langle B \rangle\rangle C$ is

defined to contain at least the points in C that satisfy B . Making these properties precise gives us the following.

$$|\text{support}(\delta)|B| \leq n \quad (59)$$

$$|\text{support}(\delta)|\bar{B}| \leq \bar{n} \quad (60)$$

$$\gamma_{\mathbb{C}}(C)|B| \subseteq \gamma_{\mathbb{C}}(\langle\langle B \rangle\rangle C) \quad (61)$$

By assumption we have the following.

$$\text{support}(\delta) \subseteq \gamma_{\mathbb{C}}(C) \quad (62)$$

$$s^{\min} \leq |\text{support}(\delta)| \leq s^{\max} \quad (63)$$

$$m^{\min} \leq \|\delta\| \leq m^{\max} \quad (64)$$

$$\forall \sigma \in \text{support}(\delta) . p^{\min} \leq \delta(\sigma) \leq p^{\max} \quad (65)$$

We need to show the following four claims.

$$\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2) \quad (66)$$

$$s_2^{\min} \leq |\text{support}(\delta_2)| \leq s_2^{\max} \quad (67)$$

$$m_2^{\min} \leq \|\delta_2\| \leq m_2^{\max} \quad (68)$$

$$\forall \sigma \in \text{support}(\delta_2) . p_2^{\min} \leq \delta_2(\sigma) \leq p_2^{\max} \quad (69)$$

Claim (66) – Support. Let $\sigma \in \text{support}(\delta_2)$. Thus it must be that $\sigma \in \text{support}(\delta)$ and $\llbracket B \rrbracket \sigma = \mathbf{true}$. By (62), we have $\sigma \in \gamma_{\mathbb{C}}(C)$, therefore $\sigma \in \gamma_{\mathbb{C}}(C_2)$ as $\{\sigma \in \gamma_{\mathbb{C}}(C) \mid \llbracket B \rrbracket \sigma = \mathbf{true}\} \subseteq \gamma_{\mathbb{C}}(C_2)$ by construction of C_2 . \square

Claim (67) – Support points. Let us write $\text{support}(\delta)$ as a union of two disjoint sets.

$$\text{support}(\delta) = \text{support}(\delta)|B| \cup \text{support}(\delta)|\bar{B}|$$

Given the disjointness of the two, we also have the following.

$$|\text{support}(\delta)| = |\text{support}(\delta)|B| + |\text{support}(\delta)|\bar{B}|$$

Now note that $\text{support}(\delta_2) = \text{support}(\delta)|B|$. Thus we can write $|\text{support}(\delta_2)| = |\text{support}(\delta)| - |\text{support}(\delta)|\bar{B}|$. We can therefore estimate the size of the support of δ_2 in the following manner.

$$\begin{aligned} |\text{support}(\delta_2)| &= |\text{support}(\delta)| - |\text{support}(\delta)|\bar{B}| \\ &\leq |\text{support}(\delta)| \\ &\leq s^{\max} \end{aligned} \quad [\text{ by (63) }]$$

Therefore, using (59) and the above, we have $|\text{support}(\delta_2)| \leq \min\{s^{\max}, n\} = s_2^{\max}$.

Going in the other direction, we can write as follows.

$$\begin{aligned} |\text{support}(\delta_2)| &= |\text{support}(\delta)| - |\text{support}(\delta)|\bar{B}| \\ &\geq s^{\min} - |\text{support}(\delta)|\bar{B}| \quad [\text{ by (63) }] \\ &\geq s^{\min} - \bar{n} \quad [\text{ by (60) }] \end{aligned}$$

Since all sets are trivially of size at least 0, we have $|\text{support}(\delta_2)| \geq \max\{s^{\min} - \bar{n}, 0\} = s_2^{\min}$. \square

Claim (69) – Probability. Note that we will show the probability claim before the mass as we will use the truth of the probability claim in the mass arguments.

Let $\sigma \in \text{support}(\delta_2)$. By definition of δ_2 , we have $\delta_2(\sigma) = \delta(\sigma)$. Thus $\sigma \in \text{support}(\delta)$ so by (65) we have:

$$p_2^{\min} = p^{\min} \leq \delta(\sigma) = \delta_2(\sigma) \leq p^{\max} = p_2^{\max} \quad \square$$

Claim (68) – Mass. Let us first show the following bound on the size of $\text{support}(\delta)|\bar{B}|$.

$$\max\{s^{\min} - n, 0\} \leq |\text{support}(\delta)|\bar{B}| \leq \min\{s^{\max}, \bar{n}\} \quad (70)$$

Since $|\text{support}(\delta)| = |\text{support}(\delta)|B| + |\text{support}(\delta)|\bar{B}|$, we can say $|\text{support}(\delta)|\bar{B}| = |\text{support}(\delta)| - |\text{support}(\delta)|B|$ and continue to the bound in the following manner.

$$\begin{aligned} |\text{support}(\delta)|\bar{B}| &= |\text{support}(\delta)| - |\text{support}(\delta)|B| \\ &\geq s^{\min} - |\text{support}(\delta)|B| \quad [\text{ by (63) }] \\ &\geq s^{\min} - n \quad [\text{ by (59) }] \end{aligned}$$

Therefore $|\text{support}(\delta)|\bar{B}| \geq \max\{s^{\min} - n, 0\}$ as claimed. For the other end of the inequality, note that we have $|\text{support}(\delta)|\bar{B}| \leq |\text{support}(\delta)| \leq s^{\max}$ by (63). Also, by (60), $|\text{support}(\delta)|\bar{B}| \leq \bar{n}$. Therefore $|\text{support}(\delta)|\bar{B}| \leq \max\{s^{\max}, \bar{n}\}$, completing our bound.

Now, let us write $\|\delta\|$ in two parts.

$$\begin{aligned} \|\delta\| &= \sum_{\sigma \in \text{support}(\delta)} \delta(\sigma) \\ &= \sum_{\sigma \in \text{support}(\delta)|B} \delta(\sigma) + \sum_{\sigma \in \text{support}(\delta)|\bar{B}} \delta(\sigma) \\ &= \|\delta_2\| + \sum_{\sigma \in \text{support}(\delta)|\bar{B}} \delta(\sigma) \end{aligned}$$

$$\text{Therefore } \|\delta_2\| = \|\delta\| - \sum_{\sigma \in \text{support}(\delta)|\bar{B}} \delta(\sigma).$$

$$\begin{aligned} \|\delta_2\| &= \|\delta\| - \sum_{\sigma \in \text{support}(\delta)|\bar{B}} \delta(\sigma) \\ &\leq m^{\max} - \sum_{\sigma \in \text{support}(\delta)|\bar{B}} \delta(\sigma) \quad [\text{ by (64) }] \\ &\leq m^{\max} - \sum_{\sigma \in \text{support}(\delta)|\bar{B}} p^{\min} \quad [\text{ by (65) }] \\ &= m^{\max} - |\text{support}(\delta)|\bar{B}| \cdot p^{\min} \\ &\leq m^{\max} - \max\{s^{\min} - n, 0\} \cdot p^{\min} \quad [\text{ by (70) }] \end{aligned}$$

Also, we can bound the mass using our other already proven conditions.

$$\begin{aligned}
\|\delta_2\| &= \sum_{\sigma \in \text{support}(\delta_2)} \delta_2(\sigma) \\
&\leq \sum_{\sigma \in \text{support}(\delta_2)} p_2^{\max} \quad [\text{by (69)}] \\
&= |\text{support}(\delta_2)| \cdot p_2^{\max} \\
&\leq s_2^{\max} \cdot p_2^{\max} \quad [\text{by (67)}]
\end{aligned}$$

Combining the bounds, we have half of our probability condition.

$$\begin{aligned}
\|\delta_2\| &\leq m_2^{\max} \\
&= \min \{ p_2^{\max} \cdot s_2^{\max}, m_2^{\max} - p^{\min} \cdot \max \{ s^{\min} - n, 0 \} \}
\end{aligned}$$

For the other half, we proceed similarly.

$$\begin{aligned}
\|\delta_2\| &= \|\delta\| - \sum_{\sigma \in \text{support}(\delta) | \bar{B}} \delta(\sigma) \\
&\geq m^{\min} - \sum_{\sigma \in \text{support}(\delta) | \bar{B}} \delta(\sigma) \quad [\text{by (64)}] \\
&\geq m^{\min} - \sum_{\sigma \in \text{support}(\delta) | \bar{B}} p^{\max} \quad [\text{by (65)}] \\
&= m^{\min} - |\text{support}(\delta) | \bar{B}| \cdot p^{\max} \\
&\geq m^{\min} - \min \{ s^{\max}, \bar{n} \} \cdot p^{\max} \quad [\text{by (70)}]
\end{aligned}$$

And likewise another bound using our other conditions.

$$\begin{aligned}
\|\delta_2\| &= \sum_{\sigma \in \text{support}(\delta_2)} \delta_2(\sigma) \\
&\geq \sum_{\sigma \in \text{support}(\delta_2)} p_2^{\min} \quad [\text{by (69)}] \\
&= |\text{support}(\delta_2)| \cdot p_2^{\min} \\
&\geq s_2^{\min} \cdot p_2^{\min} \quad [\text{by (67)}]
\end{aligned}$$

Combining the two bounds, we have the final element of our proof.

$$\begin{aligned}
\|\delta_2\| &\geq m_2^{\min} \\
&= \max \{ p_2^{\min} \cdot s_2^{\min}, m_2^{\min} - p^{\max} \cdot \min \{ s^{\max}, \bar{n} \} \}
\end{aligned}$$

F. Scalar product

Lemma 17. *If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ then $p \cdot \delta_1 \in \gamma_{\mathbb{P}}(p \cdot P_1)$.*

Proof: By assumption we have the following.

$$\begin{aligned}
\text{support}(\delta_1) &\subseteq \gamma_{\mathbb{C}}(C_1) \\
s_1^{\min} &\leq |\text{support}(\delta_1)| \leq s_1^{\max} \\
m_1^{\min} &\leq \|\delta_1\| \leq m_1^{\max} \\
\forall \sigma \in \text{support}(\delta_1) &\cdot p_1^{\min} \leq \delta_1(\sigma) \leq p_1^{\max}
\end{aligned}$$

Let $\delta_2 = p \cdot \delta_1$ and $P_2 = p \cdot P_1$. Let us assume that $p \neq 0$. In this case we need to show the following.

$$\begin{aligned}
\text{support}(\delta_1) &= \text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2) = \gamma_{\mathbb{C}}(C_1) \\
s_1^{\min} = s_2^{\min} &\leq |\text{support}(\delta_2)| = |\text{support}(\delta_1)| \leq s_2^{\max} = s_1^{\max} \\
p \cdot m_1^{\min} &= m_2^{\min} \leq \|\delta_1\| \leq m_2^{\max} = p \cdot m_1^{\max} \\
\forall \sigma \in \text{support}(\delta_2) &\cdot \\
p \cdot p_1^{\min} &= p_2^{\min} \leq \delta_2(\sigma) \leq p_2^{\max} = p \cdot p_1^{\max}
\end{aligned}$$

The first two conditions are trivially satisfied given the lack of change in the various parameters. For the mass condition, note that $\|\delta_2\| = \sum_{\sigma} \delta_2(\sigma) = \sum_{\sigma} p \cdot \delta_1(\sigma) = p \cdot \|\delta_1\|$. The probability condition is also trivially satisfied as $\delta_2(\sigma) = p \cdot \delta_1(\sigma)$.

In the case that $p = 0$, the abstract scalar product is defined with $s_2^{\min} = s_2^{\max} = p_2^{\min} = p_2^{\max} = m_2^{\min} = m_2^{\max} = 0$ and $C_2 = \emptyset_{\mathbb{C}}$. In this case note that $\text{support}(\delta_2) = \emptyset = \gamma_{\mathbb{C}}(\emptyset_{\mathbb{C}})$, and thus the conditions hold trivially. ■

G. Uniform

Lemma 62 (Soundness of Uniform). *If $\delta \in \gamma_{\mathbb{P}}(P)$ and S is uniform $x \ n_1 \ n_2$ then $\llbracket S \rrbracket \delta \in \gamma_{\mathbb{P}}(\llbracket S \rrbracket P)$.*

Proof: Recall the semantics of the statement.

$$\llbracket \text{uniform } x \ n_1 \ n_2 \rrbracket \delta = (\delta \upharpoonright_{fv(\delta)} - \{x\}) \times \delta_2$$

The distribution δ_2 is defined as follows.

$$\delta_2 = \lambda \sigma. \text{ if } n_1 \leq \sigma(x) \leq n_2 \text{ then } \frac{1}{n_2 - n_1 + 1} \text{ else } 0$$

The abstract semantics are similar.

$$\llbracket \text{uniform } x \ n_1 \ n_2 \rrbracket P = (\mathbb{f}_x(P)) \times P_2$$

Here P_2 is defined with $p_2^{\min} = p_2^{\max} = \frac{1}{n_2 - n_1 + 1}$, $s_2^{\min} = s_2^{\max} = n_2 - n_1 + 1$, $m_2^{\min} = m_2^{\max} = 1$, and $C_2 = (\{x \geq n_1, x \leq n_2\}, \{x\})$.

By construction, we have $\delta_2 \in P_2$ thus the lemma follows from Lemma 7 (Soundness of Forget) and Lemma 13 (Soundness of Product). ■

H. While loops

Definition 63. First we have some preliminary definitions. Given some set of variables, we have the following, where each distribution or state in each statement is understood to be defined over the same set of variables.

- Two distributions are *ordered*, or $\delta_1 \leq \delta_2$ iff for every state σ , $\delta_1(\sigma) \leq \delta_2(\sigma)$.
- Two probabilistic polyhedra are *ordered*, or $P_1 \sqsubseteq_{\mathbb{P}} P_2$ iff for every $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$, there exists $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$ with $\delta_1 \leq \delta_2$.
- The *zero distribution* δ is the unique distribution with $\delta(\sigma) = 0$ for every σ . We will use 0_{Dist} to refer to this distribution.

- A zero probabilistic polyhedron P , or $iszero(P_1)$ is one whose concretization contains only the zero distribution, that is $\gamma_{\mathbb{P}}(P) = \{0_{\mathbf{Dist}}\}$.

Lemma 64. *Let P_i be consistent probabilistic polyhedra, that is, $\gamma_{\mathbb{P}}(P_i) \neq \emptyset$. Then, $P_1 + P_2 \sqsubseteq_{\mathbb{P}} P_1$ iff $iszero(P_2)$.*

Proof: In the forward direction, we have $P_1 + P_2 \sqsubseteq_{\mathbb{P}} P_1$. Now, let us consider a P_2 with not $iszero(P_2)$. Thus there is $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$ with $\|\delta_2\| > 0$. Let $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ be the distribution in $\gamma_{\mathbb{P}}(P_1)$ maximizing mass, that is $\|\delta_1\| \geq \|\delta'_1\|$ for every $\delta'_1 \in \gamma_{\mathbb{P}}(P_1)$. By Lemma 12, $\delta_1 + \delta_2 \in \gamma_{\mathbb{P}}(P_1 + P_2)$ and by the definition of $P_1 + P_2 \sqsubseteq_{\mathbb{P}} P_1$, there must be $\delta_3 \in \gamma_{\mathbb{P}}(P_1)$ with $\delta_1 + \delta_2 \leq \delta_3$. Thus $\|\delta_3\| \geq \|\delta_1 + \delta_2\| = \|\delta_1\| + \|\delta_2\| > \|\delta_1\|$. This contradicts that δ_1 was mass maximizing in $\gamma_{\mathbb{P}}(P_1)$.

In the backward direction, our definition of abstract plus makes $P_1 + P_2$ identical to P_1 . Thus $P_1 + P_2 = P_1 \sqsubseteq_{\mathbb{P}} P_1$. ■

Definition 65. Given a statement $S = \text{while } B \text{ do } S'$, a distribution δ and a probabilistic polyhedron P , let us define a few useful items.

- $\omega(f) \stackrel{\text{def}}{=} \lambda\delta. f(\llbracket S' \rrbracket(\delta|B)) + \delta|\neg B$
- $\delta_1 \stackrel{\text{def}}{=} \delta$
- $\delta_{i+1} \stackrel{\text{def}}{=} \llbracket S' \rrbracket(\delta_i|B)$
- $\Delta_n \stackrel{\text{def}}{=} \sum_{i=1}^n (\delta_i|\neg B)$
- $\perp_{\mathbf{Dist}}$ is the function that takes in any distribution and produces the zero distribution $0_{\mathbf{Dist}}$, that is $\perp_{\mathbf{Dist}}(\delta) = 0_{\mathbf{Dist}}$.

Similarly we have the abstract versions of the definitions.

- $\Omega(F) \stackrel{\text{def}}{=} \lambda P. F(\langle\langle S' \rangle\rangle(P|B)) + P|\neg B$
- $P_1 \stackrel{\text{def}}{=} P$
- $P_{i+1} \stackrel{\text{def}}{=} \langle\langle S' \rangle\rangle(P_i|B)$
- $\Phi_n \stackrel{\text{def}}{=} \sum_{i=1}^n (P_i|\neg B)$
- $\perp_{\mathbb{P}}$ is a function that takes in any probabilistic polyhedron and produces a zero probabilistic polyhedron, that is $iszero(\perp_{\mathbb{P}}(P))$ for every P .

The semantics of while loops are defined as such:

$$\begin{aligned} \llbracket S \rrbracket &= \llbracket \text{while } B \text{ do } S' \rrbracket = \text{lfp}(\omega) \\ \langle\langle S \rangle\rangle &= \langle\langle \text{while } B \text{ do } S' \rangle\rangle = \text{lfp}(\Omega) \end{aligned}$$

While such definitions are of theoretical interest, they are not particularly useful for implementations, given our lack of a widening operator. Thus, our security checks will always be conditioned on termination of the abstract interpretation, defined below. We show that termination of the abstract interpretation implies termination of all corresponding concrete executions. This is crucial, as our concrete semantics (due to Clarkson et al. [9]) assumes termination to avoid leaks. To make this termination condition explicit, we provide an alternate concrete semantics for terminating while loops and show that this gives results equivalent to those of the original semantics.

Definition 66. The termination of $\llbracket S \rrbracket\delta$ is defined as follows.

- If S is an elementary statement (assignment, skip, uniform), then $\llbracket S \rrbracket\delta$ terminates.
- If S is a sequence, if statement, or a probabilistic choice statement, then $\llbracket S \rrbracket\delta$ terminates iff the various evaluations steps to evaluate S terminate. This depends on the statement type, for $S = S_1 ; S_2$, for example, it means that $\llbracket S_1 \rrbracket\delta$ terminates and so does $\llbracket S_2 \rrbracket(\llbracket S_1 \rrbracket\delta)$.
- If $S = \text{while } B \text{ do } S_1$ is a while statement, then $\llbracket S \rrbracket\delta$ terminates iff there exists n with $\delta_n = 0_{\mathbf{Dist}}$ and the evaluation steps as per definition of δ_i terminate for all i up to n .

The termination of $\langle\langle S \rangle\rangle P$ is framed similarly, except in the while case, we require the existence of n with $iszero(P_n)$ and the termination of the abstract evaluations as in the definitions of P_i for all i up to n .

The Δ_i and Φ_i capture exactly the concrete and abstract values when termination is assumed.

$$\begin{aligned} \omega^1(\perp_{\mathbf{Dist}})(\delta) &= \delta|\neg B \\ &= \Delta_1 \end{aligned}$$

$$\begin{aligned} \omega^2(\perp_{\mathbf{Dist}})(\delta) &= (\llbracket S' \rrbracket(\delta|B))|\neg B + \delta|\neg B \\ &= \delta_2|\neg B + \delta_1|\neg B \\ &= \Delta_2 \end{aligned}$$

$$\begin{aligned} \omega^i(\perp_{\mathbf{Dist}})(\delta) &= \omega^{i-1}(\perp_{\mathbf{Dist}})(\llbracket S' \rrbracket\delta|B) + \delta|\neg B \\ &= \Delta_{i-1} + \delta|\neg B \\ &= \Delta_i \end{aligned}$$

Likewise $\Omega^i(\perp_{\mathbb{P}})(P) = \Phi_i$.

Definition 67. Terminating semantics of while loops are as follows.

$$\llbracket \text{while } B \text{ do } S_1 \rrbracket\delta = \Delta_n$$

Where n is the least index with $\delta_n = 0_{\mathbf{Dist}}$. Likewise for the abstract case.

$$\langle\langle \text{while } B \text{ do } S_1 \rangle\rangle P = \Phi_n$$

Where n is the least index with $iszero(P_n)$.

Lemma 68. *If $\llbracket \text{while } B \text{ do } S_1 \rrbracket\delta$ is terminating, then $\Delta_n = (\text{lfp}(\omega))(\delta)$, noting that $\text{lfp}(\omega)$ is the original semantics of a while loop.*

Proof: As noted in [16], the evaluation of a while loop on a distribution is equal to an infinite sum:

$$\llbracket S \rrbracket\delta = \sum_{i=1}^{\infty} \delta_i|\neg B$$

By the termination assumption we have an n with $\delta_n = 0_{\mathbf{Dist}}$. Now, since $\delta_{i+1} = \llbracket S' \rrbracket\delta_i|B$ hence the mass of δ_{i+1}

cannot exceed the mass of δ_i , it is the case that if $\delta_n = 0_{\text{Dist}}$, then $\delta_i = 0_{\text{Dist}}$ for every $i \geq n$. Thus the infinite sum above can be shortened.

$$\begin{aligned} \llbracket S \rrbracket \delta &= \sum_{i=1}^{\infty} \delta_i | \neg B \\ &= \sum_{i=1}^n \delta_i | \neg B + 0_{\text{Dist}} \\ &= \Delta_n \end{aligned}$$

Remark 69 (Nature of Termination). If $\langle\langle S \rangle\rangle P$ terminates, then so must the evaluation of all of its components as defined by the semantics. This is immediate from the definition of termination.

I. Soundness of Abstraction

Theorem 6. For all P, δ , if $\delta \in \gamma_{\mathbb{P}}(P)$ and $\langle\langle S \rangle\rangle P$ terminates, then $\llbracket S \rrbracket \delta$ terminates and $\llbracket S \rrbracket \delta \in \gamma_{\mathbb{P}}(\langle\langle S \rangle\rangle P)$.

Proof: Let us show this by structural induction on S . As base cases we have the following.

- $S = \text{skip}$. In this case we have $\llbracket S \rrbracket \delta = \delta$ and $\langle\langle S \rangle\rangle P = P$. Termination is not an issue and the claim holds by assumption.
- $S = x := E$. Here non-termination is also not a possibility given non-recursive definition of assignment. Also, by Lemma 9 (Soundness of Assignment) we have $\llbracket S \rrbracket \delta \in \gamma_{\mathbb{P}}(\langle\langle S \rangle\rangle P)$.
- $S = \text{uniform } x \ n_1 \ n_2$. Again, there is no termination issues and the claim follows from Lemma 62 (Soundness of Uniform).

Let us thus assume the claim for sub-statements of S and show it for S itself. Note that the inductive assumption is general for all δ, P with $\delta \in \gamma_{\mathbb{P}}(P)$. S has several cases.

- $S = S_1 ; S_2$. By the termination remark, we know $\langle\langle S_1 \rangle\rangle P$ terminates and thus by induction $\llbracket S_1 \rrbracket \delta$ terminates and is in $\gamma_{\mathbb{P}}(\langle\langle S_1 \rangle\rangle P)$. We then apply induction once more with S_2 to find that $\llbracket S_2 \rrbracket (\llbracket S_1 \rrbracket \delta) = \llbracket S \rrbracket \delta$ terminates and is in $\gamma_{\mathbb{P}}(\langle\langle S_2 \rangle\rangle (\langle\langle S_1 \rangle\rangle P)) = \gamma_{\mathbb{P}}(\langle\langle S \rangle\rangle P)$.
- $S = \text{if } B \text{ then } S_1 \text{ else } S_2$. By the termination remark, we know that $\langle\langle S_1 \rangle\rangle (P | B)$ and $\langle\langle S_2 \rangle\rangle (P | \neg B)$ terminate. By Lemma 15 (Soundness of Conditional) we have $\delta | B \in \gamma_{\mathbb{P}}(P | B)$ and $\delta | \neg B \in \gamma_{\mathbb{P}}(P | \neg B)$. We thus apply induction to both sub-statements to conclude that $\llbracket S_1 \rrbracket (\delta | B)$ and $\llbracket S_2 \rrbracket (\delta | \neg B)$ both terminate and are in $\gamma_{\mathbb{P}}(\langle\langle S_1 \rangle\rangle (P | B))$ and $\gamma_{\mathbb{P}}(\langle\langle S_2 \rangle\rangle (P | \neg B))$ respectively. Finally we apply Lemma 12 (Soundness of Plus) to conclude $\llbracket S \rrbracket \delta = \llbracket S_1 \rrbracket (\delta | B) + \llbracket S_2 \rrbracket (\delta | \neg B) \in \gamma_{\mathbb{P}}(\langle\langle S_1 \rangle\rangle (P | B) + \langle\langle S_2 \rangle\rangle (P | \neg B)) = \gamma_{\mathbb{P}}(\langle\langle S \rangle\rangle P)$.

- $S = \text{pif } p \text{ then } S_1 \text{ else } S_2$. This case is identical to the previous except we use Lemma 17 (Soundness of Scalar Product) in place of Lemma 15 (Soundness of Conditional).
- $S = \text{while } B \text{ do } S_1$.

For this last case we must first show a claim. For every δ', P' with $\delta' \in \gamma_{\mathbb{P}}(P')$, and every i we have the following.

$$\delta'_i \in \gamma_{\mathbb{P}}(P'_i) \quad (71)$$

$$\Delta'_i \in \gamma_{\mathbb{P}}(\Phi'_i) \quad (72)$$

Let us show this claim by induction on i . As the base case we have $\delta'_1 = \delta'$ and $\Delta'_1 = \delta'_1 | \neg B = \delta' | \neg B$. Also $P'_1 = P'$ and $\Phi'_1 = P'_1 | \neg B = P' | \neg B$. By assumption we had $\delta' \in \gamma_{\mathbb{P}}(P')$ so the first part of our claim holds trivially. For the other we apply Lemma 15 (Soundness of Conditional) to conclude $\Delta'_1 \in \gamma_{\mathbb{P}}(\Phi'_1)$.

Let us assume the claim holds for all $i < n$ and show that it holds for n .

We have, by definition, $\delta'_n = \llbracket S_1 \rrbracket (\delta'_{n-1} | B)$ and $P'_n = \langle\langle S_1 \rangle\rangle (P'_{n-1} | B)$. By the (inner) induction assumption, we have $\delta'_{n-1} \in \gamma_{\mathbb{P}}(P'_{n-1})$ so by Lemma 15 we have $\delta'_{n-1} | B \in \gamma_{\mathbb{P}}(P'_{n-1} | B)$. Since $\langle\langle S \rangle\rangle P$ terminates, then so must $\langle\langle S_1 \rangle\rangle P'_{n-1} | B$ by the termination remark. Thus, by the (outer) induction hypothesis, we know that $\llbracket S_1 \rrbracket (\delta'_{n-1} | B) = \delta'_n \in \gamma_{\mathbb{P}}(\langle\langle S_1 \rangle\rangle (P'_{n-1} | B)) = \gamma_{\mathbb{P}}(P'_n)$.

For the second part of the claim, we have $\Delta'_n = \Delta'_{n-1} + \delta'_n | \neg B$ and $\Phi'_n = \Phi'_{n-1} + P'_n | \neg B$. By (inner) induction we know $\Delta'_{n-1} \in \gamma_{\mathbb{P}}(\Phi'_{n-1})$. By the first part of the claim above we know $\delta'_n \in \gamma_{\mathbb{P}}(P'_n)$ so by Lemma 15 (Soundness of Conditional) we have $\delta'_n | \neg B \in \gamma_{\mathbb{P}}(P'_n | \neg B)$. Now we apply Lemma 12 (Soundness of Plus) to conclude $\Delta'_n = \Delta'_{n-1} + \delta'_n | \neg B \in \gamma_{\mathbb{P}}(\Phi'_{n-1} + P'_n | \neg B) = \gamma_{\mathbb{P}}(\Phi'_n)$, finishing the claim.

Now, since $\langle\langle S \rangle\rangle P'$ terminates, it must be that $\langle\langle S \rangle\rangle P' = \Phi'_n$ for some n , according to the terminating semantics. Furthermore we have the following, also by definition of termination.

$$\text{iszero}(P'_n | \neg B) \quad (73)$$

This is the case since $\text{iszero}(P'_n)$ and the fact that the conditioning operation preserves $\text{iszero}(\cdot)$.

Therefore by (71) we can conclude that $\delta_n = 0_{\text{Dist}}$ as $\gamma_{\mathbb{C}}(P_n) = \{0_{\text{Dist}}\}$. Therefore $\llbracket S \rrbracket \delta$ terminates and by Lemma 68 we have $\llbracket S \rrbracket \delta = \Delta_n$. The issue of whether n is the least index with $\delta_n = 0_{\text{Dist}}$ is irrelevant as if it were not, the larger sum includes only additional 0_{Dist} terms. By (72), we have $\Delta_n \in \gamma_{\mathbb{P}}(\Phi_n)$ and we are done as $\Phi_n = \langle\langle S \rangle\rangle P$ according to the terminating semantics. ■

J. Normalization

Lemma 19. If $\delta_1 \in \gamma_{\mathbb{P}}(P_1)$ then $\text{normal}(\delta_1) \in \gamma_{\mathbb{P}}(\text{normal}(P_1))$.

Proof: By assumption we have the following.

$$\begin{aligned} \text{support}(\delta_1) &\subseteq \gamma_{\mathbb{C}}(C_1) \\ s_1^{\min} &\leq |\text{support}(\delta_1)| \leq s_1^{\max} \\ m_1^{\min} &\leq \|\delta_1\| \leq m_1^{\max} \\ \forall \sigma \in \text{support}(\delta_1) \cdot p_1^{\min} &\leq \delta_1(\sigma) \leq p_1^{\max} \end{aligned}$$

If $\|\delta_1\| = 0$ then $\text{normal}(\delta_1)$ is undefined. Since $m_1^{\min} \leq \|\delta_1\|$, it must be that $m_1^{\min} = 0$ as well, and thus $\text{normal}(P_1)$ is likewise undefined.

Let us now assume $\|\delta_1\| > 0$. Let $\delta_2 = \text{normal}(\delta_1)$ and $P_2 = \text{normal}(P_1)$. We have two sub-cases, either $m_1^{\min} = 0$ or $m_1^{\min} > 0$. In the first sub case, P_2 is defined as follows.

$$\begin{array}{l|l} p_2^{\min} = p_1^{\min}/m_1^{\max} & s_2^{\min} = s_1^{\min} \\ p_2^{\max} = 1 & s_2^{\max} = s_1^{\max} \\ m_2^{\min} = m_2^{\max} = 1 & C_2 = C_1 \end{array}$$

Since $\text{support}(\delta_2) = \text{support}(\delta_1)$, it must be that $\text{support}(\delta_2) \subseteq \gamma_{\mathbb{C}}(C_2)$ as $C_2 = C_1$. Likewise, the number of support point is unchanged in both the concrete operation and the abstract one, hence the number of support points condition for soundness are satisfied as well. Also, the probability per point in any distribution does not exceed 1 hence the p_2^{\max} condition is satisfied. As for p_2^{\min} , note that if $\sigma \in \text{support}(\delta_2) = \text{support}(\delta_1)$, we have $\delta_2(\sigma) = \delta_1(\sigma)/\|\delta_1\| \geq p_1^{\min}/\|\delta_1\| \geq p_1^{\min}/m_1^{\max}$, by assumption. Finally, $\|\delta_2\| = 1$ hence the m_2^{\min} and m_2^{\max} conditions are satisfied.

In the other case, we have $p_1^{\min} > 0$. Here P_2 is defined as follows.

$$\begin{array}{l|l} p_2^{\min} = p_1^{\min}/m_1^{\max} & s_2^{\min} = s_1^{\min} \\ p_2^{\max} = p_1^{\max}/m_1^{\min} & s_2^{\max} = s_1^{\max} \\ m_2^{\min} = m_2^{\max} = 1 & C_2 = C_1 \end{array}$$

The support, support points, total mass, and p_2^{\min} conditions are satisfied for the same reason as in the previous case. For p_2^{\max} , let $\sigma \in \text{support}(\delta_2) = \text{support}(\delta_1)$ and we have the following.

$$\begin{aligned} \delta_2(\sigma) &= \delta_1(\sigma)/\|\delta_1\| \\ &\leq p_1^{\max}/\|\delta_1\| \\ &\leq p_1^{\max}/m_1^{\min} \end{aligned}$$

K. Security

Before we prove the security theorem, let us show that the definition of abstract conditioning on a state is sound.

Lemma 70. *If $\delta \in \gamma_{\mathbb{P}}(P)$ and $\sigma_V \in \mathbf{State}_V$ with $V \subseteq \text{fv}(\delta)$ then $\delta|\sigma_V \in \gamma_{\mathbb{P}}(P|\sigma_V)$*

Proof: Recall the definition of $P|\sigma_V$.

$$P|\sigma_V = P|B$$

With $B = \bigwedge_{x \in V} (x = \sigma_V(x))$. Let us show that $\delta|\sigma_V = \delta|B$, the rest will follow from Lemma 15.

The definition of $\delta|\sigma_V$ is as follows.

$$\delta|\sigma = \lambda\sigma. \text{ if } \sigma \upharpoonright V = \sigma_V \text{ then } \delta(\sigma) \text{ else } 0$$

Meanwhile, $\delta|B$ is defined as follows.

$$\delta|B = \lambda\sigma. \text{ if } \llbracket B \rrbracket\sigma = \mathbf{true} \text{ then } \delta(\sigma) \text{ else } 0$$

The correspondence is immediate as $\llbracket B \rrbracket\sigma = \mathbf{true}$ if and only if $\sigma \upharpoonright V = \sigma_V$ as per construction of B . ■

Theorem 22. *Let δ be an attacker's initial belief. If $\delta \in \gamma_{\mathbb{P}}(P)$ and $t\text{secure}_t(S, P)$, then S is threshold secure for threshold t when evaluated with initial belief δ .*

Proof: Let us consider the contrapositive. That is, assuming $\delta \in \gamma_{\mathbb{P}}(P)$, if S is not threshold secure for t and initial belief δ , then it is not the case that $t\text{secure}_t(S, P)$.

Let $\delta_2 = \llbracket S \rrbracket\delta$ and $\delta_3 = \delta_2 \upharpoonright L$. Since S is not secure, we have $\sigma_L \in \text{support}(\delta_3)$ and $\sigma'_H \in \mathbf{State}_H$ with $(\text{normal}((\delta_2|\sigma_L) \upharpoonright H))(\sigma'_H) > t$. This implies that $(\delta_2|\sigma_L) \upharpoonright H \neq 0_{\text{Dist}}$ and therefore $\delta_2|\sigma_L \neq 0_{\text{Dist}}$ as projection preserves mass.

If $\langle\langle S \rangle\rangle P$ is not terminating, then we are done as termination is a condition for $t\text{secure}_t(S, P)$. So let us assume $\langle\langle S \rangle\rangle P$ is terminating. Let $P_2 = \langle\langle S \rangle\rangle P$. By Theorem 6, we have $\delta_2 \in \gamma_{\mathbb{P}}(P_2)$. By Lemma 70, $\delta_2|\sigma_L \in \gamma_{\mathbb{P}}(P_2|\sigma_L)$. Therefore not $\text{iszero}(P|\sigma_L)$ as $\delta_2|\sigma_L \neq 0_{\text{Dist}}$. Continuing, by Lemma 44, $(\delta_2|\sigma_L) \upharpoonright H \in \gamma_{\mathbb{P}}((P_2|\sigma_L) \upharpoonright H)$ and finally, by Lemma 19, we have $\text{normal}((\delta_2|\sigma_L) \upharpoonright H) \in \gamma_{\mathbb{P}}(\text{normal}((P_2|\sigma_L) \upharpoonright H))$. Let $\delta_4 = \text{normal}((\delta_2|\sigma_L) \upharpoonright H)$ and $P_4 = \text{normal}((P_2|\sigma_L) \upharpoonright H)$. Since $\sigma'_H \in \text{support}(\delta_4)$, we have $\delta_4(\sigma'_H) \leq p_4^{\max}$. Since $\delta_4(\sigma'_H) > t$, we have $t < p_4^{\max}$.

Also, let $P_3 = P_2 \upharpoonright L$. By Lemma 44, we have $\delta_3 \in \gamma_{\mathbb{P}}(P_3)$ so $\sigma_L \in \gamma_{\mathbb{C}}(C_3)$. We already had that not $\text{iszero}(P|\sigma_L)$ above. Thus σ_L is indeed the witness to the failure of $t\text{secure}_t(S, P_1)$. ■

APPENDIX E. SOUNDNESS PROOFS FOR $\mathcal{P}_n(\mathbb{P})$

A. Useful Lemmas

We begin with some lemmas that give properties of the concretization function for powersets of probabilistic polyhedra and addition on sets.

Lemma 71. *If $\Delta = \Delta_1 \cup \Delta_2$ then $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) = \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) + \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2)$.*

Proof: From the definition of $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ we have

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) = \sum_{P \in \Delta} \gamma_{\mathbb{P}}(P)$$

Applying $\Delta = \Delta_1 \cup \Delta_2$ and associativity of $+$ allows us to conclude

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) = \sum_{P_1 \in \Delta_1} \gamma_{\mathbb{P}}(P_1) + \sum_{P_2 \in \Delta_2} \gamma_{\mathbb{P}}(P_2)$$

Again applying the definition of $\gamma_{\mathcal{P}_n(\mathbb{P})}(\dots)$, we have

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) = \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) + \gamma_{\mathcal{P}_n(\mathbb{P})}(P_2)$$

Lemma 72. *If $D_1 \subseteq D'_1$ and $D_2 \subseteq D'_2$ then $D_1 + D_2 \subseteq D'_1 + D'_2$.*

Proof: According to the definition of addition for sets, we have

$$D_1 + D_2 = \{\delta_1 + \delta_2 \mid \delta_1 \in D_1 \wedge \delta_2 \in D_2\}$$

Consider some $\delta \in D_1 + D_2$. We have $\delta = \delta_1 + \delta_2$ with $\delta_1 \in D_1$ and $\delta_2 \in D_2$. Since $D_1 \subseteq D'_1$, we have $\delta_1 \in D'_1$. Similarly, since $D_2 \subseteq D'_2$, we have $\delta_2 \in D'_2$. Since

$$D'_1 + D'_2 = \{\delta'_1 + \delta'_2 \mid \delta'_1 \in D'_1 \wedge \delta'_2 \in D'_2\}$$

we have $\delta = \delta_1 + \delta_2 \in D'_1 + D'_2$. ■

B. Bounding Operation

Lemma 26 (Soundness of Bounding Operation). $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\lfloor \Delta \rfloor_n)$.

Proof: According to Definition 25, there are two cases for $\lfloor \Delta \rfloor_n$. If $|\Delta| \leq n$ then we have $\lfloor \Delta \rfloor_n = \Delta$ and thus $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) = \gamma_{\mathcal{P}_n(\mathbb{P})}(\lfloor \Delta \rfloor_n)$.

If $|\Delta| > n$, we reason by induction on $|\Delta|$. Since $n \geq 1$, we have that $|\Delta| \geq 2$ and thus we can partition Δ into $\Delta_1 \cup \{P_1, P_2\}$. Applying Definition 25 we then have $\lfloor \Delta \rfloor_n = \lfloor \Delta_1 \cup \{P_1 + P_2\} \rfloor_n$. The inductively-passed set has size one less than the original, allowing us to apply the inductive hypothesis to conclude the following.

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1 \cup \{P_1 + P_2\}) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\lfloor \Delta_1 \cup \{P_1 + P_2\} \rfloor_n)$$

Our conclusion will follow provided we can show

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1 \cup \{P_1 + P_2\})$$

Lemma 71 allows us to rewrite this to

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) + \gamma_{\mathcal{P}_n(\mathbb{P})}(\{P_1 + P_2\}) \quad (74)$$

We have $\Delta = \Delta_1 \cup \{P_1, P_2\}$ and thus by Lemma 71 we have

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) = \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) + \gamma_{\mathcal{P}_n(\mathbb{P})}(\{P_1, P_2\})$$

By Lemma 72, we will have (74) provided we can show

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1)$$

which is immediate, and

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\{P_1, P_2\}) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\{P_1 + P_2\})$$

The latter is proven by applying the definitions of $\gamma_{\mathcal{P}_n(\mathbb{P})}(\{P_1, P_2\})$ and $\gamma_{\mathcal{P}_n(\mathbb{P})}(\{P_1 + P_2\})$, resulting in a goal of

$$\gamma_{\mathbb{P}}(P_1) + \gamma_{\mathbb{P}}(P_2) \subseteq \gamma_{\mathbb{P}}(P_1 + P_2)$$

which follows directly from Lemma 12. ■

C. Distributive Operations

The soundness proofs for the majority of the operations on elements of $\mathcal{P}_n(\mathbb{P})$ are sound for exactly the same reason: the operations distribute over $+$, allowing us to reduce soundness for the powerset case to soundness for the case of a single probabilistic polyhedron. We start with the Lemma that is used to structure such a proof.

Lemma 73. *Consider $f : \mathbb{P} \rightarrow \mathbb{P}$, $F : \mathcal{P}_n(\mathbb{P}) \rightarrow \mathcal{P}_n(\mathbb{P})$, and $f^b : \mathbf{Dist} \rightarrow \mathbf{Dist}$. Suppose the following all hold for all δ_i, P_i .*

- 1) $f^b(\delta_1 + \dots + \delta_n) = f^b(\delta_1) + \dots + f^b(\delta_n)$
- 2) $F(\{P_1, \dots, P_n\}) = \{f(P_1), \dots, f(P_n)\}$
- 3) $\delta \in \gamma_{\mathbb{P}}(P) \Rightarrow f^b(\delta) \in \gamma_{\mathbb{P}}(f(P))$

Then $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ implies $f^b(\delta) \in \gamma_{\mathcal{P}_n(\mathbb{P})}(F(\Delta))$.

Proof: Suppose $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $\Delta = \{P_1, \dots, P_n\}$. We have the following by definition of $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$.

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta) = \gamma_{\mathbb{P}}(P_1) + \dots + \gamma_{\mathbb{P}}(P_n)$$

Applying the definition of addition on sets, we obtain

$$\gamma_{\mathbb{P}}(P_1) + \dots + \gamma_{\mathbb{P}}(P_n) = \{\delta_1 + \dots + \delta_n \mid \delta_i \in \gamma_{\mathbb{P}}(P_i)\}$$

Thus, we have that $\delta = \delta_1 + \dots + \delta_n$ where $\delta_i \in \gamma_{\mathbb{P}}(P_i)$. By premise 3 we then have $f^b(\delta_i) \in \gamma_{\mathbb{P}}(f(P_i))$ for all i .

We now consider $\gamma_{\mathcal{P}_n(\mathbb{P})}(F(\Delta))$. By premise 2 we have that this is $\gamma_{\mathcal{P}_n(\mathbb{P})}(\{f(P_1), \dots, f(P_n)\})$. Applying the definition of $\gamma_{\mathcal{P}_n(\mathbb{P})}$, this is equal to $\gamma_{\mathbb{P}}(f(P_1)) + \dots + \gamma_{\mathbb{P}}(f(P_n))$.

Expanding the definition of $+$ for sets, we have that

$$\gamma_{\mathcal{P}_n(\mathbb{P})}(F(\Delta)) = \{\delta_1 + \dots + \delta_n \mid \delta_i \in \gamma_{\mathbb{P}}(f(P_i))\}$$

Since $f^b(\delta_i) \in \gamma_{\mathbb{P}}(f(P_i))$ for all i we have $\sum_i f^b(\delta_i) \in \gamma_{\mathcal{P}_n(\mathbb{P})}(F(\Delta))$ and thus, by premise 1 we have $f^b(\sum_i \delta_i) \in \gamma_{\mathcal{P}_n(\mathbb{P})}(F(\Delta))$ and thus $f^b(\delta) \in \gamma_{\mathcal{P}_n(\mathbb{P})}(F(\Delta))$ as desired. ■

Lemma 74 (Soundness of Forget). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ then $f_y(\delta) \in \gamma_{\mathcal{P}_n(\mathbb{P})}(f_y(\Delta))$.*

Proof: We will apply Lemma 73 with $f^b = \lambda \delta. \delta \upharpoonright (f_v(\delta) - \{y\})$, $f = \lambda P. f_y(P)$, and $F = \lambda \Delta. f_y(\Delta)$. Lemma 7 gives us premise 3. The definition of $f_y(\Delta)$ satisfies premise 2. Let $V = f_v(\delta) - \{y\}$. It remains to show premise 1, which states

$$(\delta_1 + \dots + \delta_n) \upharpoonright V = \delta_1 \upharpoonright V + \dots + \delta_n \upharpoonright V$$

We show this for the binary case, from which the n -ary version above follows.

$$(\delta_1 + \delta_2) \upharpoonright V = \delta_1 \upharpoonright V + \delta_2 \upharpoonright V$$

Expanding the definition of projection, we then obtain the following goal.

$$\begin{aligned} \lambda\sigma_V \in \mathbf{State}_V. \quad & \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} (\delta_1 + \delta_2)(\sigma') = \\ \lambda\sigma_V \in \mathbf{State}_V. \quad & \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_1(\sigma') \\ & + \lambda\sigma_V \in \mathbf{State}_V. \quad \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_2(\sigma') \end{aligned}$$

We can now apply the definition of $+$ for distributions to the right-hand side to obtain a goal of

$$\begin{aligned} \lambda\sigma_V \in \mathbf{State}_V. \quad & \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} (\delta_1 + \delta_2)(\sigma') = \\ \lambda\sigma_V \in \mathbf{State}_V. \quad & \left(\sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_1(\sigma') + \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_2(\sigma') \right) \end{aligned}$$

These functions are equal if they give equal results for all inputs. Thus, we must show the following for all σ_V .

$$\begin{aligned} \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} (\delta_1 + \delta_2)(\sigma') = \\ \left(\sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_1(\sigma') + \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_2(\sigma') \right) \end{aligned}$$

Finally, applying the definition of $+$ for distributions to the left-hand side of the equality yields

$$\begin{aligned} \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} (\delta_1(\sigma') + \delta_2(\sigma')) = \\ \left(\sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_1(\sigma') + \sum_{\sigma' | (\sigma' \upharpoonright V = \sigma_V)} \delta_2(\sigma') \right) \end{aligned}$$

This follows by associativity and commutativity of $+$. \blacksquare

Lemma 75 (Soundness of Projection). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $V \subseteq \text{fv}(\delta)$ then $\delta \upharpoonright V \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta \upharpoonright V)$.*

Proof: Inductive application of Lemma 74 (Soundness of Forget) as was the case in the base domain. \blacksquare

Lemma 76 (Soundness of Assignment). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ then $\delta[x \rightarrow E] \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta[x \rightarrow E])$.*

Proof: As in Lemma 74, we apply Lemma 73. We have premises 3 (by Lemma 9) and 2 (by definition) and must show premise 1. This means showing that

$$(\delta_1 + \delta_2)[x \rightarrow E] = \delta_1[x \rightarrow E] + \delta_2[x \rightarrow E]$$

Expanding the definition of assignment, we must show that the following

$$\lambda\sigma. \quad \sum_{\tau | \tau[x \rightarrow [E]]\tau = \sigma} (\delta_1 + \delta_2)(\tau)$$

is equal to

$$\left(\lambda\sigma. \quad \sum_{\tau | \tau[x \rightarrow [E]]\tau = \sigma} \delta_1(\tau) \right) + \left(\lambda\sigma. \quad \sum_{\tau | \tau[x \rightarrow [E]]\tau = \sigma} \delta_2(\tau) \right)$$

Again applying the definition of $+$ for distributions and using extensional equality for functions yields the following goal, which follows by associativity and commutativity of $+$.

$$\begin{aligned} \forall\sigma. \quad & \left(\sum_{\tau | \tau[x \rightarrow [E]]\tau = \sigma} (\delta_1(\tau) + \delta_2(\tau)) = \right. \\ & \left. \sum_{\tau | \tau[x \rightarrow [E]]\tau = \sigma} \delta_1(\tau) + \sum_{\tau | \tau[x \rightarrow [E]]\tau = \sigma} \delta_2(\tau) \right) \end{aligned}$$

\blacksquare

Lemma 77 (Soundness of Scalar Product). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ then $p \cdot \delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(p \cdot \Delta)$.*

Proof: This proof follows the same format as the others in this section. We apply Lemma 9 with the definition of scalar product for powersets and Lemma 17. We must show

$$p \cdot (\delta_1 + \delta_2) = p \cdot \delta_1 + p \cdot \delta_2$$

Expanding according to the definition of scalar product and $+$ for distributions, we obtain the following as a goal.

$$\lambda\sigma. p \cdot (\delta_1(\sigma) + \delta_2(\sigma)) = \lambda\sigma. p \cdot \delta_1(\sigma) + p \cdot \delta_2(\sigma)$$

The result follows by distributivity of \cdot over $+$. \blacksquare

Lemma 78 (Soundness of Conditioning). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ then $\delta|B \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta|B)$.*

Proof: Again we apply Lemma 9, this time using Lemma 15 to satisfy premise 3. We let $f^b = \lambda\delta. \delta|B$, $f = \lambda P. P|B$, and $F = \lambda\Delta. \Delta|B$. We must show

$$(\delta_1 + \delta_2)|B = \delta_1|B + \delta_2|B$$

Applying the definition of conditioning and addition for distributions, we have to show the following for all σ .

$$\begin{aligned} \text{if } [B]\sigma \text{ then } (\delta_1 + \delta_2)(\sigma) \text{ else } 0 = \\ (\text{if } [B]\sigma \text{ then } \delta_1(\sigma) \text{ else } 0) + \\ (\text{if } [B]\sigma \text{ then } \delta_2(\sigma) \text{ else } 0) \end{aligned}$$

We proceed via case analysis. If $[B]\sigma = \mathbf{false}$ then we have $0 = 0 + 0$, which is a tautology. If $[B]\sigma = \mathbf{true}$, we have to show

$$(\delta_1 + \delta_2)(\sigma) = \delta_1(\sigma) + \delta_2(\sigma)$$

which follows directly from the definition of $+$ on distributions. \blacksquare

D. Other Powerset Lemmas

We now show the lemmas for operations in the powerset domain that do not immediately follow from distributivity over plus of the operations in the base domain.

Lemma 79 (Soundness of Product). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $\delta' \in \gamma_{\mathcal{P}_n(\mathbb{P})}(P')$ and $\text{fv}(\Delta) \cap \text{fv}(P') = \emptyset$ then $\delta \times \delta' \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta \times P')$.*

Proof: Let $\Delta = \{P_1, \dots, P_n\}$. We first expand definitions in our goal, obtaining

$$\delta \times \delta' \in \gamma_{\mathbb{P}}(P_1 \times P') + \dots + \gamma_{\mathbb{P}}(P_n \times P')$$

Applying the definition of addition for sets, we obtain a goal of

$$\delta \times \delta' \in \left\{ \sum_i \delta_i \mid \delta_i \in \gamma_{\mathbb{P}}(P_i \times P') \right\}$$

This holds provided we can find $\delta_i \in \gamma_{\mathbb{P}}(P_i \times P')$ such that $\delta \times \delta' = \sum_i \delta_i$. We have from $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ that $\delta = \sum_j \delta_j$ for some $\delta_j \in \gamma_{\mathbb{P}}(P_j)$. We then have from Lemma 13 and $\delta' \in \gamma_{\mathcal{P}_n(\mathbb{P})}(P')$ and $\text{fv}(\Delta) \cap \text{fv}(P') = \emptyset$ that $\delta_j \times \delta' \in \gamma_{\mathbb{P}}(P_j \times P')$ for all j . We now show that the δ_i we were searching for are these $\delta_j \times \delta'$. To do so, we must show that $\delta \times \delta' = \sum_j (\delta_j \times \delta')$. We have $\delta = \sum_j \delta_j$ and thus the result follows by distributivity of \times over $+$, which we show now.

Goal: \times distributes over $+$: We want to show the following when $\text{domain}(\delta_1) = \text{domain}(\delta_2)$ and $\text{domain}(\delta_1) \cap \text{domain}(\delta') = \emptyset$.

$$(\delta_1 + \delta_2) \times \delta' = \delta_1 \times \delta' + \delta_2 \times \delta'$$

Expanding the definition of $+$ and of \times , we obtain

$$\begin{aligned} \lambda(\sigma, \sigma'). (\delta_1(\sigma) + \delta_2(\sigma)) \cdot \delta'(\sigma') = \\ \lambda(\sigma, \sigma'). (\delta_1(\sigma) \cdot \delta'(\sigma') + \delta_2(\sigma) \cdot \delta'(\sigma')) \end{aligned}$$

This holds due to distributivity of \cdot over $+$. \blacksquare

Lemma 80 (Soundness of Addition). *If $\delta_1 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1)$ and $\delta_2 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2)$ then $\delta_1 + \delta_2 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1 + \Delta_2)$.*

Proof: First let us take care of the special cases that occur when $\text{iszero}(\Delta_1)$ or $\text{iszero}(\Delta_2)$. Without the loss of generality let us say $\text{iszero}(\Delta_2)$. The sum is defined to be identical to Δ_1 . Since $\text{iszero}(\Delta_2)$, it must be that $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2)$ contains only the zero distribution 0_{Dist} , therefore $\delta_2 = 0_{\text{Dist}}$. Therefore $\delta_1 + \delta_2 = \delta_1$ and by assumption, $\delta_1 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) = \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1 + \Delta_2)$.

In the case where Δ_1 and Δ_2 are both non-zero, we have $\Delta_1 + \Delta_2 = \lfloor \Delta_1 \cup \Delta_2 \rfloor_n$. Suppose $\delta_1 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1)$ and $\delta_2 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2)$. By Lemma 71 we have $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) + \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2) = \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1 \cup \Delta_2)$. The set $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) + \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2)$ is $\{\delta'_1 + \delta'_2 \mid \delta'_1 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1) \wedge \delta'_2 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2)\}$. Our distributions δ_1 and δ_2 satisfy these conditions and thus are in $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1 \cup \Delta_2)$. It remains to show that $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_1 \cup \Delta_2) \subseteq \gamma_{\mathcal{P}_n(\mathbb{P})}(\lfloor \Delta_1 \cup \Delta_2 \rfloor_n)$, but this is exactly Lemma 26. \blacksquare

E. Main Soundness Theorem for Powerset Domain

The main soundness theorem is an identical restatement of the main soundness theorem in the base domain and the proof is likewise identical, save for replacement of the relevant base domain definitions and lemmas with the powerset ones. The only corresponding lemma which has not yet been proven follows below.

Lemma 81. *Let Δ_i be consistent probabilistic polyhedron sets, that is, $\gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_i) \neq \emptyset$. Then, $\Delta_1 + \Delta_2 \sqsubseteq_{\mathbb{P}} \Delta_1$ iff $\text{iszero}(\Delta_2)$.*

Proof: The proof is identical to the Lemma 64, replacing the base domain lemmas and definitions with the powerset ones. \blacksquare

Theorem 24 (Soundness of Abstraction). *For all δ, S, Δ , if $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $\langle\langle S \rangle\rangle \Delta$ terminates, then $\llbracket S \rrbracket \delta$ terminates and $\llbracket S \rrbracket \delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\langle\langle S \rangle\rangle \Delta)$.*

Proof: The proof is identical to the main soundness proof for the base domain (Theorem 6), replacing definitions and lemmas about the base domain abstraction with the corresponding definitions and lemmas about the powerset domain. \blacksquare

Lemma 82 (Soundness of Normalization). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ then $\text{normal}(\delta) \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\text{normal}(\Delta))$.*

Proof: Whenever $\|\delta\| = 0$, the normalization in the concrete sense is undefined, likewise it is undefined in the abstract sense. So let us assume $\|\delta\| > 0$.

Let $\underline{m} = \sum_i m_i^{\min}$ and $\overline{m} = \sum_i m_i^{\max}$. By assumption we have $\delta = \sum_i \delta_i$ with

$$\delta_i \in \gamma_{\mathbb{P}}(P_i) \quad (75)$$

Thus we have $\|\delta\| = \sum_i \|\delta_i\|$ and we conclude $\underline{m} = \sum_i m_i^{\min} \leq \|\delta\| \leq \sum_i m_i^{\max} = \overline{m}$ via (75).

$$\underline{m} \leq \|\delta\| \leq \overline{m} \quad (76)$$

Let $\delta' = \text{normal}(\delta) = \frac{1}{\|\delta\|} \delta = \sum_i \frac{1}{\|\delta\|} \delta_i$, due to linearity of scalar product. Let us thus show that $\frac{1}{\|\delta\|} \delta_i \in \gamma_{\mathbb{P}}(\text{normal}(P_i)(\underline{m}, \overline{m})) = \gamma_{\mathbb{P}}(\text{normal}(\Delta))$ which would conclude the proof. Let us write $P_{i'} = \text{normal}(P_i)(\underline{m}, \overline{m})$ and $\delta_{i'} = \frac{1}{\|\delta\|} \delta_i$. We must thus show the following.

$$\text{support}(\delta_{i'}) \subseteq \gamma_{\mathbb{C}}(C_{i'}) \quad (77)$$

$$s_{i'}^{\min} \leq |\text{support}(\delta_{i'})| \leq s_{i'}^{\max} \quad (78)$$

$$m_{i'}^{\min} \leq \|\delta_{i'}\| \leq m_{i'}^{\max} \quad (79)$$

$$\forall \sigma \in \text{support}(\delta_{i'}) . p_{i'}^{\min} \leq \delta_{i'}(\sigma) \leq p_{i'}^{\max} \quad (80)$$

Claim (77) holds trivially as $\text{support}(\delta_{i'}) = \text{support}(\delta_i)$, $C_{i'} = C_i$, and (75). Claim (78) holds due to the same reasoning.

For (79), in the case where $\underline{m} > 0$, we reason, via (76), as follows.

$$\begin{aligned} m_i^{\min} &\leq \|\delta_i\| \leq m_i^{\max} \\ \frac{m_i^{\min}}{\underline{m}} &\leq \frac{1}{\|\delta\|} \|\delta_i\| \leq \frac{m_i^{\max}}{\underline{m}} \\ \frac{m_i^{\min}}{\underline{m}} &\leq \left\| \frac{1}{\|\delta\|} \delta_i \right\| \leq \frac{m_i^{\max}}{\underline{m}} \\ m_{i'}^{\min} &= \frac{m_i^{\min}}{\underline{m}} \leq \|\delta_{i'}\| \leq \frac{m_i^{\max}}{\underline{m}} = m_{i'}^{\max} \end{aligned}$$

If $\underline{m} = 0$, the definition of normalization makes $m_{i'}^{\max} = 1$, which is also sound as all distributions have mass no more than 1.

The (80) claim is shown using reasoning identical to the mass claim above. \blacksquare

Lemma 28 (Soundness of Simple Maximal Bound Estimate). *If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\{P_i\})$ and $P = \sum_i P_i$ then $\max_{\sigma} \delta(\sigma) \leq p^{\max}$.*

Proof: By assumption we have $\delta = \sum_i \delta_i$ with $\delta_i \in \gamma_{\mathbb{P}}(P_i)$ thus by Lemma 12 (Soundness of Plus), we have $\delta \in \gamma_{\mathbb{P}}(\sum_i P_i) = \gamma_{\mathbb{P}}(P)$, thus for every $\sigma \in \text{support}(\delta)$, $\delta(\sigma) \leq p^{\max}$, hence $\max_{\sigma} \delta(\sigma) \leq p^{\max}$. \blacksquare

The above lemma shows soundness of the very simple method of estimating the maximum probability but in the implementation we use the method based on poly partitioning and the following lemma.

Lemma 32. $\max_{pp}(\Delta) \stackrel{\text{def}}{=} \max_{\sigma \in R} \Delta^{\max}(\sigma) = \max_{\sigma} \Delta^{\max}(\sigma)$ where \mathcal{L} is a poly partition of Δ and R is a representative set of \mathcal{L} .

Proof: Let \mathcal{L} be the poly partition of $\Delta = \{C_i\}$ as in the statement of the lemma. Let us first show a claim: if $\sigma, \sigma' \in L \in \mathcal{L}$ then

$$A \stackrel{\text{def}}{=} \{C \in \Delta \mid \sigma \in \gamma_{\mathcal{C}}(C)\} = \{C \in \Delta \mid \sigma' \in \gamma_{\mathcal{C}}(C)\} \stackrel{\text{def}}{=} B \quad (81)$$

Let $C \in A$. Thus $\sigma \in \gamma_{\mathcal{C}}(C)$ so by Definition 31 (2), we have $\sigma \in \gamma_{\mathcal{C}}(L')$ for some $L' \in \mathcal{L}$. By (1) it must be that $L = L'$ and by (3), we have $\gamma_{\mathcal{C}}(L) = \gamma_{\mathcal{C}}(L') \subseteq \gamma_{\mathcal{C}}(C)$. Therefore $\sigma' \in \gamma_{\mathcal{C}}(C)$ and thus $C \in B$, showing $A \subseteq B$. The other direction is identical, concluding $A = B$ as claimed. \square

Now we can get back to the main lemma. Let σ^* be the state with $\Delta^{\max}(\sigma^*) = \max_{\sigma} \Delta^{\max}(\sigma)$. Thus $\sigma^* \in \gamma_{\mathcal{C}}(L)$ for some $L \in \mathcal{L}$, by Definition 31 (2). Let σ_L be any representative of L , that is $\sigma_L \in \gamma_{\mathcal{C}}(L)$.

$$\begin{aligned} \Delta^{\max}(\sigma^*) &= \sum_i P_i^{\max}(\sigma^*) \\ &= \sum_{i \mid \sigma^* \in \gamma_{\mathcal{C}}(C_i)} p_i^{\max} \\ &= \sum_{i \mid \sigma_L \in \gamma_{\mathcal{C}}(C_i)} p_i^{\max} \quad [\text{by (81)}] \\ &= \sum_i P_i^{\max}(\sigma_L) \\ &= \Delta^{\max}(\sigma_L) \end{aligned}$$

Now we see that $\max_{\sigma} \Delta^{\max}(\sigma) = \Delta^{\max}(\sigma^*) = \Delta^{\max}(\sigma_L) = \max_{pp}(\Delta)$ as claimed. \blacksquare

Before we prove the security theorem, let us show that the definition of abstract conditioning on a state is sound.

Lemma 83. *If $\delta \in \gamma_{\mathbb{P}}(\Delta)$ and $\sigma_V \in \text{State}_V$ with $V \subseteq \text{fv}(\delta)$ then $\delta \mid \sigma_V \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta \mid \sigma_V)$*

Proof: Recall the definition of $\Delta \mid \sigma_V$.

$$\Delta \mid \sigma_V = \Delta \mid B$$

With $B = \bigwedge_{x \in V} (x = \sigma_V(x))$. Let us show that $\delta \mid \sigma_V = \delta \mid B$, the rest will follow from Lemma 78.

The definition of $\delta \mid \sigma_V$ is as follows.

$$\delta \mid \sigma = \lambda \sigma. \text{ if } \sigma \upharpoonright V = \sigma_V \text{ then } \delta(\sigma) \text{ else } 0$$

Meanwhile, $\delta \mid B$ is defined as follows.

$$\delta \mid B = \lambda \sigma. \text{ if } \llbracket B \rrbracket \sigma = \text{true} \text{ then } \delta(\sigma) \text{ else } 0$$

The correspondence is immediate as $\llbracket B \rrbracket \sigma = \text{true}$ if and only if $\sigma \upharpoonright V = \sigma_V$ as per construction of B . \blacksquare

Theorem 35 (Soundness for Threshold Security). *Let δ be an attacker's initial belief. If $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$ and $t\text{secure}_t(S, \Delta)$, then S is threshold secure for threshold t when evaluated with initial belief δ .*

Proof: Let us consider the contrapositive. That is, assuming $\delta \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta)$, if S is not threshold secure for t and initial belief δ , then it is not the case that $t\text{secure}_t(S, \Delta)$.

Let $\delta_2 = \llbracket S \rrbracket \delta$ and $\delta_3 = \delta_2 \upharpoonright L$. Since S is not secure, we have $\sigma_L \in \text{support}(\delta_3)$ and $\sigma'_H \in \text{State}_H$ with $(\text{normal}((\delta_2 \mid \sigma_L) \upharpoonright H))(\sigma'_H) > t$. This implies that $(\delta_2 \mid \sigma_L) \upharpoonright H \neq 0_{\text{Dist}}$ and therefore $\delta_2 \mid \sigma_L \neq 0_{\text{Dist}}$ as projection preserves mass.

If $\llbracket S \rrbracket \Delta$ is not terminating, then we are done as termination is a condition for $t\text{secure}_t(S, \Delta)$. So let us assume $\llbracket S \rrbracket \Delta$ is terminating. Let $\Delta_2 = \llbracket S \rrbracket \Delta$. By Theorem 24, we have $\delta_2 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2)$. By Lemma 83, $\delta_2 \mid \sigma_L \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_2 \mid \sigma_L)$. Therefore not $\text{iszero}(\Delta_2 \mid \sigma_L)$. Continuing, by Lemma 75, $(\delta_2 \mid \sigma_L) \upharpoonright H \in \gamma_{\mathcal{P}_n(\mathbb{P})}((\Delta_2 \mid \sigma_L) \upharpoonright H)$ and finally, by Lemma 82, we have $\delta_4 \stackrel{\text{def}}{=} \text{normal}((\delta_2 \mid \sigma_L) \upharpoonright H) \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\text{normal}((\Delta_2 \mid \sigma_L) \upharpoonright H))$. Let $\Delta_4 = \text{normal}((\Delta_2 \mid \sigma_L) \upharpoonright H)$.

By Remark 30, we have $\delta_4(\sigma'_H) \leq \max_{\sigma} \Delta_4^{\max}(\sigma)$ and by Lemma 32 we have $\max_{\sigma} \Delta_4^{\max}(\sigma) = \max_{pp}(\Delta_4)$. But $\delta_4(\sigma'_H) > t$ so $\max_{pp}(\Delta_4) > t$, a potential failure of $t\text{secure}_t(S, \Delta)$.

To finish the proof we need to make sure that σ_L was indeed a valid witness to the failure of $t\text{secure}_t(S, P_1)$. Let

$\Delta_3 = \{P_i''\} = \Delta_2 \upharpoonright L$. By Lemma 75, we have $\delta_3 \in \gamma_{\mathcal{P}_n(\mathbb{P})}(\Delta_3)$ so $\delta_3 = \sum_i \delta'_i$ with $\delta'_i \in \gamma_{\mathbb{P}}(P_i'')$. Since $\sigma_L \in \text{support}(\delta_3)$ it must be that $\delta_3(\sigma_L) > 0$ and thus $\delta'_i(\sigma_L) > 0$ for at least one i . Thus $\sigma_L \in \text{support}(\delta'_i) \subseteq \gamma_{\mathbb{C}}(C_i'')$ for at

least one i and therefore $\sigma_L \in \gamma_{\mathcal{P}(\mathbb{C})}(\{C_i''\})$. Also, we have already shown that not $\text{iszero}(\Delta_2 \upharpoonright \sigma_L)$, thus σ_L is indeed the witness as needed. ■