Universal equivalence and majority of probabilistic programs over finite fields

Gilles Barthe

MPI-SP & IMDEA Software Institute

Charlie Jacomme

LSV, CNRS & ENS Paris-Saclay &
Inria & Université Paris-Saclay

Steve Kremer

LORIA, Inria Nancy-Grand Est &
CNRS & Université de Lorraine

Abstract

We study decidability problems for equivalence of probabilistic programs, for a core probabilistic programming language over finite fields of fixed characteristic. The programming language supports uniform sampling, addition, multiplication and conditionals and thus is sufficiently expressive to encode boolean and arithmetic circuits. We consider two variants of equivalence: the first one considers an interpretation over the finite field $\mathbb{F}_q$, while the second one, which we call universal equivalence, verifies equivalence over all extensions $\mathbb{F}_{q^k}$ of $\mathbb{F}_q$. The universal variant typically arises in provable cryptography when one wishes to prove equivalence for any length of bitstrings, i.e., elements of $\mathbb{F}_{2^k}$ for any k . While the first problem is obviously decidable, we establish its exact complexity which lies in the counting hierarchy. To show decidability, and a doubly exponential upper bound, of the universal variant we rely on results from algorithmic number theory and the possibility to compare local zeta functions associated to given polynomials. Finally we study several variants of the equivalence problem, including a problem we call majority, motivated by differential privacy.

CCS Concepts: • Security and privacy → Logic and verification; • Theory of computation → Program reasoning.

Keywords: program equivalence, probabilistic programs, finite fields, decidability and complexity

ACM Reference Format:

Gilles Barthe, Charlie Jacomme, and Steve Kremer. 2020. Universal equivalence and majority of probabilistic programs over finite fields. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '20)*, July 8–11, 2020, Saarbrücken, Germany.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
LICS '20, July 8–11, 2020, Saarbrücken, Germany
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7104-9/20/07...\$15.00

<https://doi.org/10.1145/3373718.3394746>

ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3373718.3394746>

1 Introduction

Program equivalence is one of the most fundamental tools in the theory of programming languages and arguably the most important example of relational property. Program equivalence has been studied extensively, leading to numerous decidability results and sound proof methods. This paper is concerned with the decidability of equivalence and relational properties for a core imperative probabilistic programming language. Like many other probabilistic programming languages, our language supports sampling from distributions, and conditioning distributions on an event. The specificity of our language is that it operates over finite fields of the form $\mathbb{F}_{q^k}$. Therefore, expressions are interpreted as polynomials and assertions are boolean combinations of polynomial identities. Sampling is interpreted using the uniform distributions over sets defined by assertions, and branching and conditioning are relative to assertions.

We consider two relational properties, equivalence and majority, which we define below, and several related properties, which we explain in the next paragraph. For each property, we consider two variants of the problem. In the first variant, which we call the fixed case, the value of k is fixed. In the second variant, which we call the universal variant, we require the property to hold for all possible values of k . Consider two programs P_1 and P_2 with m inputs and n outputs. These programs are interpreted as functions $\llbracket P_1 \rrbracket^{q^k}, \llbracket P_2 \rrbracket^{q^k} : \mathbb{F}_{q^k}^m \rightarrow \text{Distr}(\mathbb{F}_{q^k}^n)$.

- q^k -equivalence (denoted $P_1 \approx_{q^k} P_2$) requires that P_1 and P_2 define the same distributions: $\llbracket P_1 \rrbracket^{q^k} = \llbracket P_2 \rrbracket^{q^k}$. Equivalently, for every input $\vec{a} \in \mathbb{F}_{q^k}^m$ and output $\vec{b} \in \mathbb{F}_{q^k}^n$,

$$\mathbb{P}\{\vec{x} = \vec{b} \mid \vec{x} \leftarrow \llbracket P_1 \rrbracket^{q^k}(\vec{a})\} = \mathbb{P}\{x = \vec{b} \mid \vec{x} \leftarrow \llbracket P_2 \rrbracket^{q^k}(\vec{a})\}.$$

q^∞ -equivalence requires the property to hold on all extensions of a field, i.e.,

$$P_1 \approx_{q^\infty} P_2 \text{ iff } \forall k. P_1 \approx_{q^k} P_2$$

- q^k -majority requires that for a fixed $r \in \mathbb{Q}$, and for every input $\vec{a} \in \mathbb{F}_{q^k}^m$ and output $\vec{b} \in \mathbb{F}_{q^k}^n$, we have

$$\mathbb{P}\{\vec{x} = \vec{b} \mid \vec{x} \stackrel{\$}{\leftarrow} \llbracket P_1 \rrbracket^{q^k}(\vec{a})\} \leq r \cdot \mathbb{P}\{\vec{x} = \vec{b} \mid \vec{x} \stackrel{\$}{\leftarrow} \llbracket P_2 \rrbracket^{q^k}(\vec{a})\}.$$

q^k -0-majority (denoted $P_1 <_{q^k}^r P_2$) is a variant of majority, where we only consider the output $b = 0^n$, rather than quantifying over all outputs. q^∞ -0-majority requires the property to hold on all extensions of a field, i.e.,

$$P_1 <_{q^\infty}^r P_2 \text{ iff } \forall k. P_1 <_{q^k}^r P_2$$

The following two boolean programs illustrate the difference between equivalence and universal equivalence.

Example 1.1.

$$x \stackrel{\$}{\leftarrow} \mathbb{F}; \text{ return } (x^2 + x) \quad \Bigg| \quad \text{return } 0$$

are 2- but not 2^2 -equivalent, and hence not 2^∞ -equivalent. Indeed, when instantiating $\mathbb{F}$ with $\mathbb{F}_2$, the left hand side program simply evaluates to zero, which is not the case with $\mathbb{F}_4$. On the other hand, the programs

$$x \stackrel{\$}{\leftarrow} \mathbb{F}; \text{ return } (x) \quad \Bigg| \quad x \stackrel{\$}{\leftarrow} \mathbb{F}; \text{ return } (x + 1)$$

are q^∞ -equivalent as both programs define the uniform distribution over $\mathbb{F}$, whatever finite field is used for the interpretation of $\mathbb{F}$. These examples also illustrate the difference with the well-studied polynomial identity testing (PIT) problem, as the first two programs are 2-equivalent, while PIT does not consider $x^2 + x$ and 0 to be equal on $\mathbb{F}_2$, nor would Q_1 and Q_2 be considered identical.

The fixed and universal variants of the equivalence and majority problems are directly inspired from applications in security and privacy. In the fixed setting, the equivalence and majority problems are related to probabilistic non-interference and differential privacy. The relationships between probabilistic non-interference and equivalence and between differential privacy and majority are explained informally as follows:

- probabilistic non-interference: for simplicity, assume that P has two inputs x (secret) and y (public), and a single (public) output. For every x , let P_x be the unique program such that $P_x(y) = P(x, y)$. Then P is non-interfering iff for every x_1 and x_2 , the two programs P_{x_1} and P_{x_2} are equivalent.
- differential privacy: for simplicity consider the case where the base field is $\mathbb{F}_2$. For every program P with n inputs, define the residual programs $P_{i,0}$ and $P_{i,1}$ obtained by fixing the i -th output to 0 and 1 respectively. Then the program P is $\log(r)$ -differentially private iff for every i , $P_{i,0}$ and $P_{i,1}$ (and $P_{i,1}$ and $P_{i,0}$) satisfy r -majority.

In the universal setting, the parameter k can loosely be understood as the security parameter. Universal equivalence is

a special case of statistical indistinguishability and as such arises naturally in provable security, where the goal is to prove (depending on applications either as end goal, or as an intermediate goal) that two programs are equivalent for all possible interpretations (e.g. for all possible lengths of bitstrings, i.e. for all $\mathbb{F}_{2^k}$).

Summary of results

We also consider the following problems, which are also motivated by security and privacy and are directly related to equivalence:

- (bounded) simulatability: given programs P_1 and P_2 , does there exist a context $C[\cdot]$ (of bounded degree) such that $C[P_1]$ is equivalent to P_2 ;
- independence: are outputs Y and Y' of program P independent conditioned on Z , i.e. for every input x , is the distribution of Y independent from the distribution of Y' , when conditioning on the value of Z ? Although independence is not naturally expressed as a relational property, it has been shown in [4] that relational methods are useful for proving independence.

The first contribution of the paper is a systematic study of the complexity of the aforementioned problems in the fixed setting. We prove that the q^k -equivalence problem is $\text{coNP}^{C=P}$ -complete for any fixed k . We also study the special case of *linear* programs, i.e. multiplication, conditional and conditioning free, for which the problem can be decided in polynomial time. For the majority problem, we consider two settings: programs with and without inputs. We show that the k -majority problem for inputless programs is PP-complete, whereas the k -majority for arbitrary programs is coNP^{PP} -complete—thus the second problem is strictly harder than the first, unless $\text{PH} \subset \text{PP}^1$. The proofs are given by reductions to MAJSAT and E-MAJSAT respectively. Note that we do not include any result about bounded simulatability in the finite case, since we only derive easy consequences of equivalence. These results complement recent work on the complexity of checking differential privacy for arithmetic circuits [15], see Related Work below.

The second, and main contribution, is the study of universal equivalence, q^∞ -equivalence for short, and universal (0-)majority, q^∞ -(0-)majority for short. First, we show that the q^∞ -equivalence problem is in 2-EXP and $\text{coNP}^{C=P}$ -hard.

Our proof is based on local zeta Riemann functions, a powerful tool from algebraic geometry, that characterize the number of zeros of a tuple of polynomials in all extensions of a finite field. Lauder and Wan [19] notably propose an algorithm to compute such functions, whose complexity is however exponential. Based on this result, our proof proceeds in three steps.

¹As $\text{PH} \subset \text{coNP}^{\text{PP}}$, $\text{PP} = \text{coNP}^{\text{PP}}$ would imply $\text{PH} \subset \text{PP}$ which is commonly believed to be false.

First, we give a reduction for arithmetic programs (no conditionals, nor conditioning) from universal equivalence to checking that some specific local zeta Riemann functions are always null. Then, we reduce the general case to programs without conditioning, and programs without conditioning to arithmetic programs. To justify the use of the local zeta Riemann functions, we also provide counterexamples why simpler methods fail or only provide sufficient conditions. Our decidability result significantly generalize prior work on universal equivalence [3], which considers the case of linear programs, see Related Work below. In the special case of *arithmetic* programs, i.e., programs without conditionals nor conditioning, equivalence can be decided in EXP-time, rather than 2-EXP.

Second, we give an exponential reduction from the universal 0-majority problem to the positivity problem for Linear Recurrence Sequences (LRS), which given a LRS, asks whether it is always positive. Despite its apparent simplicity, the positivity problem remains open. Decidability has been obtained independently by Mignotte et al [25] and by Vereshchagin [32] for LRS of order ≤ 4 and later by Ouaknine and Worrell [29] for LRS with order ≤ 5 . Moreover, Ouaknine and Worrell prove in the same paper that deciding positivity for LRS of order 6 would allow to solve hard open problems in Diophantine approximation. In the general case, the best known lower bound for the positivity problem is NP-hardness [28].

Our reduction is based on the observation that the Taylor series of any rational function satisfies a linear recurrence sequence. Therefore, every tuple of polynomials yields a linear recurrence sequence via its local zeta Riemann function. Unfortunately, the order of the linear recurrence sequence is related to the degree of the local zeta Riemann function, and thus decidability results for small orders do not apply. This suggests that the problem may not have an efficient solution. Using the results from [18], we observe that the reduction extends to a more general form of universal majority problem.

Finally, we obtain lower complexity bounds by reducing the finite case to the universal case. It remains an interesting open question whether the universal case is strictly harder than the finite case.

Figures 1 and 2 summarize our results for the equivalence and majority problems.

Related work

Universal equivalence. The case of linear programs is studied in [3]. The authors propose a decision procedure for universal equivalence based on the classic XOR-lemma [11]. We give an alternative decision procedure and analyze its complexity.

The case of linear programs with random oracles is considered in [9]. The authors give a polynomial time decision

procedure for computational indistinguishability of two inputless programs. Informally, computational indistinguishability is an approximate notion of universal equivalence, stating that the statistical distance between the output of two programs on the same input is upper bounded by a negligible function of the parameter k . Their proof is based on linear algebra.

The case of pseudo-linear (i.e. linear with conditionals) programs is considered in [17]. The authors consider the universal simulatability problem, rather than the universal equivalence problem. The crux of their analysis is a completeness theorem for pseudo-linear functions. In Section 4.3, we show that universal equivalence reduces to universal simulatability. As [17] shows the decidability of universal simulatability for pseudo-linear programs, it therefore follows that universal equivalence of pseudo-linear programs is decidable.

Fixed equivalence. There is a vast amount of literature on proving equivalence of probabilistic programs. We only review the most relevant work here.

Murawski and Ouaknine [26] prove decidability of equivalence of second-order terms in probabilistic ALGOL. Their proof is based on a fully abstract game semantics and a connection between program equivalence and equivalence of probabilistic automata.

Legay et al [20] prove decidability of equivalence for a probabilistic programming language over finite sets. Their language supports sampling from non-uniform distributions, loops, procedure calls, and open code, but not conditioning. They show that program equivalence can be reduced to language equivalence for probabilistic automata, which can be decided in polynomial time.

Barthe et al [5] develop a relational program logic for probabilistic programs without conditioning. Their logic has been used extensively for proving program equivalence, with applications in provable security and side-channel analysis.

Majority problems. The closest related work develops methods for proving differential privacy or for quantifying information flow.

Frederikson and Jha [14] develop an abstract decision procedure for satisfiability modulo counting, and then use a concrete instantiation of their procedure for checking representative examples from multi-party computation.

Barthe et al [2] show decidability of ϵ -differential privacy for a restricted class of programs. They allow loops and sampling from Laplace distributions, but impose several other constraints on programs. An important aspect of their work is that programs are parametrized by $\epsilon > 0$, so their decision procedure establishes ϵ -differential privacy for all values of ϵ . Technically, their decision procedure relies on the decidability of a fragment of the reals with exponentials by McCallum and Weispfenning [24].

	x -(conditional)-{equivalence, independence, uniformity}			q^∞ -simulatability
	linear	arithmetic	general	
$x = q^k$	PTIME	$\text{coNP}^{C=P}$ -complete	$\text{coNP}^{C=P}$ -complete	decidable
$x = q^\infty$	PTIME	EXP $\text{coNP}^{C=P}$ -hard	2-EXP $\text{coNP}^{C=P}$ -hard	$\text{coNP}^{C=P}$ -hard

Figure 1. Summary of results related to equivalence

	q^k -0-majority	q^k -majority	q^∞ -0-majority	q^∞ -majority
without inputs	PP-complete	coNP^{PP} -complete	PP-hard	
			$\leq_{\text{EXP}}$ POSITIVITY	
with inputs	coNP^{PP} -complete			?
			coNP^{PP} -hard	

Figure 2. Summary of results related to majority

Gaboardi, Nissim and Purser [15] study the complexity of verifying pure and approximate (ϵ, δ) -differential privacy for arithmetic programs, as well as approximations of the parameters ϵ and δ . The parameter δ quantifies the approximation and $\delta = 0$ corresponds to the pure case. Our majority problem can be seen as a subcase of differential privacy, where r corresponds to ϵ , and $\delta = 0$. In particular, the complexity class they obtain for pure differential privacy coincides with the complexity of our 0-majority problem, even when restricted to the case $r = 1$. This means that the ϵ parameter does not essentially contribute to the complexity of the verification problem. Also, while they consider arithmetic programs, we consider the more general case of programs with conditioning.

Chistikov, Murawski and Purser [10] also study the complexity of approximating differential privacy, but in the case of Markov Chains.

Theory of fields. A celebrated result by Ax [1] shows that the theory of finite fields is decidable. In a recent development based on Ax's result, Johnson [16] proves decidability of the theory of rings extended with quantifiers $\mu_k^n x. P$, stating that the number of x such that P holds is equal to k modulo n . Although closely related, these results do not immediately apply to the problem of equivalence.

2 Programming Language

We consider a high-level probabilistic programming language with sampling from semi-algebraic sets and conditioning, as well as a more pure, yet equi-expressive, core language that can encode all previous constructs and define its formal semantics.

$P ::=$	<i>polynomials</i>
$i \in \mathbb{F}_q$	fixed value
x	variable
$P_1 + P_2$	field addition
$P_1 \times P_2$	field multiplication
$b ::=$	<i>boolean conditions</i>
$P = 0$	atomic formula
$b_1 \wedge b_2$	and
$b_1 \vee b_2$	or
$\neg b$	not
$e ::=$	<i>program expressions</i>
$x := P$	assignment
$\vec{r} \xleftarrow{\$} \{X \in \mathbb{F}^m \mid b\}$	sampling
observe b	observe
$e_1; e_2$	sequential composition
if b then e_1 else e_2	conditional branching
return $(P_1, \dots, P_n)$	return of arity n

Figure 3. Program syntax

2.1 Syntax and informal semantics

We define in Figure 3 the syntax for simple probabilistic programs (without loops nor recursion²). Our programs will operate on finite fields. We denote by $\mathbb{F}_q$ the (unique) finite field with q elements, where $q = p^s$ for some integer s and prime p . Programs are parametrized by a finite field $\mathbb{F}$, which will be instantiated by some $\mathbb{F}_{q^k}$ during the interpretation. Given a polynomial $P \in \mathbb{F}_q[x_1, \dots, x_m]$ and $X \in \mathbb{F}_{q^k}^m$, we denote by $P(X)$ the evaluation of P given X inside $\mathbb{F}_{q^k}$.

²Universal equivalence for programs over finite fields with loops becomes undecidable.

The expressions of our programs provide constructs for assigning a polynomial P to a variable ($x := P$), as well as, for randomly sampling values. With for instance $\vec{r} = r_1, \dots, r_m$, the expression $r_1, \dots, r_m \stackrel{\$}{\leftarrow} \{X \in \mathbb{F}^m \mid b\}$ uniformly samples m values from the set of m -tuples of values in $\mathbb{F}$ such that the condition b holds, and assigns them to variables $r_1, \dots, r_m$. For example, $r \stackrel{\$}{\leftarrow} \{x \in \mathbb{F} \mid 0 = 0\}$ (which we often simply write $r \stackrel{\$}{\leftarrow} \mathbb{F}$) uniformly samples a random element in $\mathbb{F}$, while $r_1, r_2 \stackrel{\$}{\leftarrow} \{x_1, x_2 \in \mathbb{F}^2 \mid \neg(x_1 = 0)\}$ samples two random variables, ensuring that the first one is not 0. Note that the use of polynomial conditions allows to express any rational distribution over the base field $\mathbb{F}_q$.

The construct `observe  $b$` allows to condition the continuation by b : if b evaluates to false the program fails; the semantics of a program is the conditional distribution where b holds. Expressions also allow classical constructs for sequential composition, conditional branching and returning a result.

In a well-formed program we suppose that every variable is bound at most once, and if it is bound, then it is only used after the binding. Unbound variables correspond to the inputs of the program. We moreover suppose that each branch of a program P ends with a `return` instruction that returns the same number n of elements; n is then called the arity of the program and denoted $|P|$. Given two sets of variables I and R , we denote by $\mathcal{P}_q(I, R)$ the set of such well-formed programs, where I is the set of unbound variables (intuitively, the set of input variables) and R the set of variables sampled by the program.

Example 2.1. Consider the following simple program :

```
if  $i = 0$  then return 0
else  $r \stackrel{\$}{\leftarrow} \mathbb{F}$ ; observe  $r \times i = 1$ ; return  $r$ 
```

This program defines a probabilistic algorithm for computing the inverse of a field element i . If i is 0, by convention the algorithm returns 0. Otherwise, the algorithm uniformly samples an element r . This is obviously not a practical procedure for computing an inverse, but we use it to illustrate the semantics of conditioning. The `observe` instruction checks whether r is the inverse of i . If this is the case we return r , otherwise the program fails. As we will see below, our semantics normalizes the probability distribution to only account for non-failing executions. Hence, this algorithm will return the inverse of any positive i with probability 1. Equivalently, this program can be written by directly conditioning the sample, replacing “ $r \stackrel{\$}{\leftarrow} \mathbb{F}$; observe $r \times i = 1$,” by “ $r \stackrel{\$}{\leftarrow} \{x \in \mathbb{F} \mid x \times i = 1\}$,”.

2.2 A core language

While the above introduced syntax is convenient for writing programs, we introduce a more pure, core language that is

actually equally expressive and will ease the technical developments in the remaining of the paper. To define this core language, we add an explicit failure instruction $\perp$, similarly to [7]. It allows us to get rid of conditioning in random samples and observe instructions. Looking ahead, and denoting by $\llbracket P \rrbracket^{q^k}$ the semantics of the program P inside $\mathbb{F}_{q^k}$, we will have that $\llbracket r_1, \dots, r_m \stackrel{\$}{\leftarrow} \{X \in \mathbb{F}^m \mid b\}; e \rrbracket^{q^k} =$

$$\llbracket r_1, \dots, r_m \stackrel{\$}{\leftarrow} \mathbb{F}^m; \text{if } b \text{ then } e \text{ else } \perp \rrbracket^{q^k}$$

and $\llbracket \text{observe } b; e \rrbracket^{q^k} = \llbracket \text{if } b \text{ then } e \text{ else } \perp \rrbracket^{q^k}$. Without loss of generality, we can inline deterministic assignments, and use code motion to perform all samplings eagerly, i.e., all random samplings are performed upfront. Therefore we can simply consider that each variable in R is implicitly uniformly sampled in $\mathbb{F}_{q^k}$. Programs are then tuples of simplified expressions $(e_1, \dots, e_n)$ defined as follows.

$e ::=$	simplified expressions
$ P$	polynomial
$ \perp$	failure
$ \text{if } b \text{ then } e_1 \text{ else } e_2$	conditional branching

We suppose that all nested tuples are flattened and write (P, Q) to denote the program which simply concatenates the outputs of P and Q . When clear from the context, we may also simply write $\vec{0}$ instead of the all zero tuple $(0, \dots, 0)$. We denote by $\bar{\mathcal{P}}_q(I, R)$ the set of arithmetic programs, that are simply tuples of polynomials. Remark that arithmetic programs cannot fail.

One may note that the translation from the surface language to the core language is not polynomial in general. Indeed, constructs of the form `(if  $b$  then  $x := t_1$  else  $x := t_2$ ;  $P$ )`, i.e. sequential composition after a conditional, implies to propagate the branching over the assignment to all branches of P , and doubles the number of conditional branchings of P . All complexity results will be given for the size of the program given inside the core language. Remark that in a functional style version of the surface language, where we replace $x := t$ by `let  $x = t$  in` and removed sequential composition, the translation would however be polynomial. Similarly, for the class of programs without sequential composition after conditional branchings, the translation is also polynomial.

2.3 Semantics

We now define the semantics of our core language. The precise translation from the high level syntax previously presented and our core language is standard and omitted.

Deterministic semantics. We first define a *deterministic* semantics where all random samplings have already been defined. For a set X of variables, with $P \in \mathbb{F}_q[X]$ and $\vec{x} \in \mathbb{F}_{q^k}^{|X|}$, $P(\vec{x})$ classically denotes the evaluation of P inside $\mathbb{F}_{q^k}$. We also denote $b(\vec{v})$ the evaluation of a boolean test, where all polynomials are evaluated according to $\vec{v}$. For a program

$e \in \mathcal{P}_q(I, R)$ and $\vec{v} \in \mathbb{F}_{q^k}^{|I \cup R|}$, we define a natural evaluation of e , denoted $[e]_{\vec{v}}^{q^k}$, which is a value inside $\mathbb{F}_{q^k}^{|P|} \times \{\perp\}$:

$$\begin{aligned} [P]_{\vec{v}}^{q^k} &= P(\vec{v}) \text{ where } P \in \mathbb{F}_q[I \uplus R] \\ [\perp]_{\vec{v}}^{q^k} &= \perp \\ \left[\begin{array}{l} \text{if } b \text{ then } e_1 \\ \text{else } e_2 \end{array} \right]_{\vec{v}}^{q^k} &= \begin{cases} [e_1]_{\vec{v}}^{q^k} & \text{if } b(\vec{v}) \text{ holds on } \mathbb{F}_{q^k} \\ [e_2]_{\vec{v}}^{q^k} & \text{else} \end{cases} \\ [(e_1, \dots, e_n)]_{\vec{v}}^{q^k} &= \begin{cases} \perp & \text{if } [e_i]_{\vec{v}}^{q^k} = \perp \text{ for some } i \\ ([e_1]_{\vec{v}}^{q^k}, \dots, [e_n]_{\vec{v}}^{q^k}) & \text{else} \end{cases} \end{aligned}$$

Intuitively, the set of executions corresponding to non failure executions represent the set of possible executions of the program. We next define probabilistic semantics by sampling uniformly the valuations of the random variables while conditioning on the fact that the program does not fail.

Probabilistic semantics. For any n , the set of distributions over $\mathbb{F}_{q^k}^n$ is denoted by $\text{Distr}(\mathbb{F}_{q^k}^n)$. For a program $P \in \mathcal{P}_q(I, R)$ with $|P| = n$, and $|I| = m$, we define its semantics to be a function from inputs to a distribution over the outputs:

$$\llbracket P \rrbracket^{q^k} : \mathbb{F}_{q^k}^m \rightarrow \text{Distr}(\mathbb{F}_{q^k}^n)$$

We assume that programs inside $P \in \mathcal{P}_q(I, R)$ do not fail all the time, i.e., for any possible input and any program its probability of failure is strictly less than 1. For program P , input $\vec{i} \in \mathbb{F}_{q^k}^m$ and output $\vec{o} \in \mathbb{F}_{q^k}^n$ we set $\mathbb{P}\{\vec{x} = \vec{o} \mid \vec{x} \leftarrow \llbracket P \rrbracket^{q^k}(\vec{i})\}$ to

$$\frac{\mathbb{P}\{[P]_{\vec{i}, \vec{r}}^{q^k} = \vec{o} \mid \vec{r} \xleftarrow{\$} \mathbb{F}_{q^k}^{|R|}\}}{\mathbb{P}\{[P]_{\vec{i}, \vec{r}}^{q^k} \neq \perp \mid \vec{r} \xleftarrow{\$} \mathbb{F}_{q^k}^{|R|}\}}$$

Note that the normalization by conditioning on non-failing programs is well defined as we supposed that programs do not always fail.

3 The fixed case

We start by studying the complexity of several problems over a given finite field. In this case we only manipulate finite objects, and hence all problems are obviously decidable, by explicitly computing the distributions. We however provide precise complexity results and show that these problems have complexities in the counting hierarchy [31]. We also define the universal variant and state some results that are common to both variants of the problems.

3.1 Conditional equivalence

In this section, we prove that for any $k \in \mathbb{N}$, the q^k -equivalence problem is $\text{coNP}^{\text{C=P}}$ -complete. To this end, we introduce a technical generalization of the equivalence problem, that we

call q^k -conditional equivalence, and we proceed in four steps, showing that:

1. without loss of generality, we can consider programs without inputs; (Lemma 3.2)
2. verifying if the conditioned distributions of two inputless programs coincide on a fixed point is in C=P ; (Lemma 3.3)
3. verifying if the conditioned distribution of inputless programs coincide on all points is in $\text{coNP}^{\text{C=P}}$; (Corollary 3.4)
4. and finally, even equivalence for programs over $\mathbb{F}_2$ is $\text{coNP}^{\text{C=P}}$ -hard. (Lemma 3.5)

3.1.1 Defining conditional equivalence. q^k -conditional equivalence is a generalization of equivalence, where we require programs to be equivalent when the distributions are conditioned by some other program being equal to zero. Conditional equivalence is a technical generalisation, that is interesting because it is self-reducible when removing for instance the conditionals.

Definition 3.1 (q^k -conditional equivalence). Let $P_1, Q_1 \in \mathcal{P}_q(I, R)$ and $P_2, Q_2 \in \overline{\mathcal{P}}_q(I, R)$ with $|P_1| = |Q_1| = n$. We denote $P_1 \mid P_2 \approx_{q^k} Q_1 \mid Q_2$, if:

$$\forall \vec{i} \in \mathbb{F}_{q^k}^{|I|}. \forall \vec{c} \in \mathbb{F}_{q^k}^n. \llbracket (P_1, P_2) \rrbracket_{\vec{i}}^{q^k}(\vec{c}, \vec{0}) = \llbracket (Q_1, Q_2) \rrbracket_{\vec{i}}^{q^k}(\vec{c}, \vec{0})$$

The universal version q^∞ -conditional equivalence is defined similarly to q^∞ -equivalence. Note that conditional equivalence is a direct generalization of equivalence, as for $P, Q \in \mathcal{P}_q(I, R)$, $P \approx_{q^k} Q$ if and only if $P \mid 0 \approx_{q^k} Q \mid 0$.

We also remark that equivalence over $\mathbb{Z}$ is undecidable, which is a consequence of Hilbert's 10th problem, as a polynomial over randomly sampled variables will be equivalent to zero if and only if it does not have any solutions.

We first define precisely the decision problems associated to our questions, for $k \in \mathbb{N} \cup \{\infty\}$:

q^k -conditional equivalence
INPUT: $P_1, Q_1 \in \mathcal{P}_q(I, R)$, $P_2, Q_2 \in \overline{\mathcal{P}}_q(I, R)$. $P_1 \mid P_2 \approx_{q^k} Q_1 \mid Q_2$?

The decision problem for q^k -equivalence simply corresponds to q^k -conditional equivalence with P_2 and Q_2 being equal to 0. In the following we will show that both problems are interreducible, and that q^k -equivalence and q^k -conditional equivalence are both $\text{coNP}^{\text{C=P}}$ -complete.

3.1.2 Complexity results. Recall that C=P -complete is the set of decision problems solvable by a NP Turing Machine whose number of accepting paths is equal to the number of rejecting paths. halfSAT is the natural C=P -complete problem: is a CNF boolean formula ϕ satisfied for exactly half of its valuations. $\text{coNP}^{\text{C=P}}$ is the set of decision problems whose complement can be solved by a NP Turing Machine with

access to an oracle deciding problems in $C=P$. The canonical $\text{coNP}^{C=P}$ problem is (using the results from [30, Sec. 4] and [22]) A-halfSAT: given a CNF boolean formula $\phi(X, Y)$, for all valuations of X is the formula satisfied for exactly half of the valuations of Y .

Also, recall that conditional equivalence is a direct generalization of equivalence. We thus trivially have, for any $k \in \mathbb{N} \cup \{\infty\}$, that q^k -equivalence reduces in polynomial time to q^k -conditional equivalence.

We first study the complexity of deciding if the distributions of two programs are equal on a specific point. To do so, we remark that it is not necessary to consider inputs when considering equivalence or conditional equivalence. The intuition is that inputs can be seen as random values, that must be synchronized on both sides. This synchronization is achieved by explicitly adding these random variables to the output, forcing them to have the same value on both side. The following Lemma is a generalization to conditional equivalence of a Lemma from [4].

Lemma 3.2. *For any $k \in \mathbb{N} \cup \{\infty\}$, q^k -conditional equivalence reduces to q^k -conditional equivalence restricted to programs without inputs in polynomial time.*

Omitted proofs can be found in [6]. As we can without loss of generality ignore the inputs, we study the complexity of deciding if the distributions of two inputless programs coincide on a specific point. To this end, we build a Turing Machine, such that it will accept half of the time if and only if the programs given as input have the same probability to be equal to some given value. Essentially, it is based on the fact that over $\mathbb{F}_2$, if $r = 0$ then $P \text{ else } (Q + 1) \approx_2 r$ if and only if $P \approx_2 Q$.

Lemma 3.3. *Let $P_1, Q_1 \in \mathcal{P}_q(\emptyset, R)$ and $P_2, Q_2 \in \overline{\mathcal{P}}_q(\emptyset, R)$ with $|P_1| = |Q_1| = n$. For any $\vec{c} \in \mathbb{F}_{q^k}^n$, we can decide in $C=P$ if:*

$$\llbracket (P_1, P_2) \rrbracket^{q^k}(\vec{c}, \vec{0}) = \llbracket Q_1, Q_2 \rrbracket^{q^k}(\vec{c}, \vec{0})$$

As $C=P$ is closed under finite intersection [30], we can decide in $C=P$ if two distributions over a set of fixed size are equal, by testing the equality over all points. When we only consider inputless programs of fixed arity, the set of points to test is constant, and the equivalence problem is in $C=P$ (see [6] for details). However, when we extend to inputs, or to programs of variable arity, we need to be able to check for all possible points if the distribution are equal over this point. (Note that our encoding that allows to only consider inputless programs increases the arity.) Checking all possible points is typically in coNP . We thus obtain that:

Corollary 3.4. *q^k -equivalence and q^k -conditional equivalence are in $\text{coNP}^{C=P}$ for any $k \in \mathbb{N}$.*

To conclude completeness for both q^k -equivalence and q^k -conditional equivalence, it is sufficient to show the hardness of 2-equivalence, which we do by reducing A-halfSAT.

We simply transform a CNF boolean formula into a polynomial over $\mathbb{F}_2$, such that the polynomial is uniform if the formula is in A-halfSAT. This is a purely technical operation (see [6]).

Lemma 3.5. *2-equivalence is $\text{coNP}^{C=P}$ -hard.*

The results about q^k -conditional equivalence naturally translate to the independence problem: are the distribution of multiple programs independent? We obtain the same complexity class by proving that the problems are interreducible. These results are detailed in [6].

3.2 Majority

The goal of this section is to show that the majority problem is coNP^{PP} -complete. To this end, we study the complexity of q^k -0-majority, showing:

- PP-completeness for inputless programs;
- coNP^{PP} -completeness in general.

The proof in both cases uses similar ideas as for equivalence. Note that we actually use the same Turing Machine for the Membership. As both complexity classes are closed under finite intersection, it yields the complexity of q^k -majority, which can be decided using q^k times q^k -0-majority.

3.2.1 The majority problem. q^k -majority asks if, given two programs, the quotient of their distribution is bounded on all points by some rational r . q^k -0-majority is a subcase, where we only ask if the quotient of their distribution is bounded on a single point. This problem allows to estimate the distance between two distributions. It is close to the differential privacy question, which asks, when $\delta = 0$, if the quotient of two distributions is bounded over all points by some e^ϵ .

We observe that the majority problem is harder than equivalence, as majority for $r = 1$ implies equivalence. An important difference between equivalence and majority is that the presence of inputs actually changes the complexity of the majority problem.

Let us define the decision problem associated to q^k -majority, with $k \in \mathbb{N} \cup \{\infty\}$:

q^k -majority
INPUT: $P, Q \in \mathcal{P}_q(I, R), r \in \mathbb{Q}$.
$P \prec_{q^k}^r Q$?

We consider that r is given in input as two integers written in unary. Essentially, this is because if one wishes to encode any r , it requires an exponential blow up, but in practice, we tend to use some particular rationals such as $r = q^l$, for which there is no exponential blow up.

3.2.2 Complexity results. We recall that PP is the set of languages accepted by a polynomial time non-deterministic Turing Machine where the acceptance condition is that a majority of paths are accepting. A natural PP-complete problem

is MAJSAT: is a boolean CNF formula satisfied for at least half of its valuations. coNP^{PP} is the class of problems whose complement is decided by a NP Turing Machine with access to an oracle deciding problems in PP. The classical NP^{PP} problem is E-MAJSAT [22], which given a CNF boolean formula $\phi(X, Y)$, asks if there is a valuation of X such that $\phi(X, Y)$ is satisfied for at least half of the valuation of Y .

The complexity of q^k -0-majority over inputless programs is derived in a similar way as for equivalence. (j'ai remplacé un paragraphe par ça) The only difficulty is that we are comparing with a rational. We thus briefly show how one can assume without loss of generality that $r = 1$ (in which case we omit r from the notation). The idea is, given $r, s \in \mathbb{N}$, that $P <_{q^k}^{\frac{r}{s}} Q \Leftrightarrow (P, T_r) <_{q^k} (Q, T_s)$, if T_j is a machine which is equal to zero with probability $\frac{1}{j}$.

Lemma 3.6. *For any $k \in \mathbb{N}$, q^k -0-majority reduces in polynomial time to q^k -0-majority with $r = 1$.*

The proof showing that q^k -0-majority is in PP is similar to proving that testing if two distributions are equal over a point is in C=P . We prove PP-completeness by deriving the hardness from MAJSAT.

Lemma 3.7. *For any $k \in \mathbb{N}$, q^k -0-majority restricted to inputless programs is PP-complete.*

Finally, as PP is closed under finite intersection, we also get that q^k -majority over inputless programs with a fixed arity is PP-complete.

Let us now turn to the general version, for programs with inputs. By using some fresh inputs variables, let us remark that one can easily reduce q^k -majority to q^k -0-majority. Indeed, for $P, Q \in \mathcal{P}_q(I, R)$ and $c \in \mathbb{F}_q^{|I|}$, with a fresh $x \in I$:

$$\begin{aligned} \forall i \in \mathbb{F}_q^{|I|}. \llbracket P \rrbracket_i^{q^k}(c) &\leq r \llbracket Q \rrbracket_i^{q^k}(c) \\ &\Leftrightarrow \\ (P - x) &<_{q^k}^r (Q - x) \end{aligned}$$

We show that q^k -majority is coNP^{PP} complete, and thus is most likely³ harder than its version without inputs. The membership and hardness proofs are similar to the equivalence problem when going from C=P to $\text{coNP}^{\text{C=P}}$.

Lemma 3.8. *q^k -majority is coNP^{PP} complete.*

4 The universal case

In this section we first give some general insights on universal equivalence showing important differences with the case of a fixed field. Then we provide our main decidability result, first for arithmetic programs, then arithmetic programs enriched with conditionals, and finally for general programs. We continue by studying two other problems in the universal case: simulatability and 0-majority.

³As $\text{PH} \subset \text{coNP}^{\text{PP}}$, $\text{PP} = \text{coNP}^{\text{PP}}$ would imply $\text{PH} \subset \text{PP}$ which is commonly believed to be false.

4.1 General remarks

In this section we try to provide some insights on the difficulty of deciding q^∞ -equivalence. First of all, we note that equivalence and universal equivalence do *not* coincide.

Example 4.1. The program $x^2 + x$ and the program 0 are equivalent over $\mathbb{F}_2$ (they are then both equal to zero), but not over $\mathbb{F}_4$.

In the case of a given finite field, equivalence can be characterized by the existence of a bijection, see for instance [4]. We denote by $\text{bij}_{\mathbb{F}_q}^m$ the set of bijections over $\mathbb{F}_q^m$. Any element $\sigma \in \text{bij}_{\mathbb{F}_q}^m$ can be expressed as a tuple of polynomials (see e.g. [27]), and can be applied as a substitution. The characterization can then be stated as follows, where we denote by $=_{\mathbb{F}_q}$ equality between polynomials modulo the rule of the field (i.e., $X^q = X$).

$$P \approx_q Q \Leftrightarrow \exists \sigma \in \text{bij}_{\mathbb{F}_q}^m, P =_{\mathbb{F}_q} Q \sigma$$

However, there are universally equivalent programs such that there does *not* exist a universal σ suitable for all extensions.

Example 4.2. Consider, $P = xy + yx + zx$, with $\sigma : (x, y, z) \mapsto (x, y + x, z + x)$, we get that $P \approx_{2^\infty} x^2 + yz$. Now, $x \mapsto x^2$ is a bijection over all $\mathbb{F}_{2^k}$, so we also have $P \approx_{2^\infty} x + yz$ and finally $P \approx_{2^\infty} x$.

But here, a bijection between $x^2 + yz$ and x must use the inverse of x^2 whose expression depends on the size of the field. Thus, there isn't a universal polynomial σ which is a bijection such that on all $\mathbb{F}_{2^k}$, $P =_{\mathbb{F}_{2^k}} Q \circ \sigma$.

Nevertheless, we can note that for linear programs this characterization allows us to show that q -equivalence and q^∞ -equivalence are equivalent. Intuitively, the bijection allowing to obtain the equality between two linear programs is also a bijection valid for all extensions of the finite field, as the bijection is linear, and is thus a witness of equivalence over all extensions. For linear programs, there exists a polynomial time decision procedure for equivalence, and hence for universal equivalence.

Lemma 4.3. *q^∞ -equivalence restricted to linear programs is in PTIME.*

Moreover, building on results from [23] on Tame automorphisms, we can use the above characterization to design a sufficient condition which implies universal equivalence for general programs. Even though not complete this sufficient condition may be useful to verify universal equivalence more efficiently in practice. These results are detailed in [6].

4.2 Decidability of universal equivalence

We show decidability of q^∞ -equivalence, leveraging tools from algebraic geometry, showing that:

1. q^∞ -conditional equivalence is decidable for arithmetic programs; (Lemma 4.5)

2. it is also decidable for programs with conditionals;
([Lemma 4.7](#))
3. it is finally decidable for programs with conditioning,
e.g. failures. ([Lemma 4.8](#))

We first recall the definition and relevant properties of local zeta Riemann functions. For a tuple P of polynomials $P_1, \dots, P_m \in \mathbb{F}_q[X_1, \dots, X_n]$, the local zeta Riemann function over T is the formal series

$$Z(P, T) = \exp \left(\sum_{k \in \mathbb{N}^*} \frac{|N_k(P)|}{k} T^k \right)$$

where $N_k(P) = \{\vec{x} \in \mathbb{F}_{q^k}^n \mid \bigwedge_{1 \leq i \leq m} P_i(\vec{x}) = 0\}$. Weil's conjecture [33] states several fundamental properties of local zeta Riemann functions over algebraic varieties. Dwork [12] proves part of Weil's conjecture stating that the local zeta Riemann functions over algebraic varieties is a rational function with integer coefficients—recall that $Z(T)$ is a rational function iff there exist polynomials $R(T)$ and $S(T)$ such that $Z(T) = R(T)/S(T)$. Bombieri [8] shows that the sum of the degrees of R and S is upper bounded by $4(d+9)^{n+1}$, where d is the total degree of $(P_1, \dots, P_m)$. It follows that the values of N_k for $k \leq 4(d+9)^{n+1}$ suffice for computing Z ; since these values can be computed by brute force, this yields an algorithm for computing Z . We will by abuse of notations write $Z(P)$ instead of $Z(P, T)$ for the local zeta function of P . $Z(P)$ completely characterizes the number of times P is equal to zero on all the different extensions. For instance, $Z(P) = Z(Q)$ allows us to conclude that P and Q always evaluate to zero for the same number of valuations, and this over any $\mathbb{F}_{q^k}$. As Z can effectively be computed [19], we can use it to decide q^∞ -equivalence.

Notice that, given two programs P and Q , the local zeta function directly allow us to conclude if they are equal to some value with the same probability for all extensions of the base field. Moreover, thanks to [18], the computability of the local zeta function can be extended from counting the number of points such that $P = 0$ for a tuple of polynomials, to counting the number of points such that ϕ holds, where ϕ is an arbitrary first order formula over finite fields.

Corollary 4.4. *Let ϕ and ψ be two first order formulae built over atoms of the form $P = 0$ with $P \in \mathbb{F}_q[X]$, and with free variables $F \subset X$. One can decide if for all $k \in \mathbb{N}$:*

$$\left| \{\vec{f} \in \mathbb{F}_{q^k}^{|F|} \mid \phi(\vec{f}) = 1\} \right| = \left| \{\vec{f} \in \mathbb{F}_{q^k}^{|F|} \mid \psi(\vec{f}) = 1\} \right|$$

Thus, for any two events which can be expressed as a first order formula over finite field one can verify if they happen with the same probability over all extensions of the base field. Remark that this cannot be used to decide universal equivalence, as equivalence cannot be expressed inside a first order formula.

We first show that, thanks to the local zeta functions, q^∞ -equivalence is decidable for arithmetic programs, i.e programs without conditionals or conditioning.

Lemma 4.5. *Let $P_1, P_2, Q_1, Q_2 \in \overline{\mathcal{P}}_q(\emptyset, R)$.*

$$\begin{aligned} P_1 \mid P_2 & \quad Z((Q_1 - Q_1\sigma, Q_2, Q_2\sigma)) \\ \approx_{q^\infty} & \Leftrightarrow Z((P_1 - Q_1\sigma, P_2, Q_2\sigma)) \\ Q_1 \mid Q_2 & \quad Z((P_1 - P_1\sigma, P_2, P_2\sigma)) \end{aligned}$$

where $\sigma : R \mapsto R'$ maps each variable to a fresh one.

Proof. We assimilate $P_1P_2, Q_1, Q_2 \in \overline{\mathcal{P}}_q(\emptyset, R)$ of size m with polynomials, denoting $P(X)$ the value of P given $X \in \mathbb{F}_{q^k}^m$. Given an enumeration $1 \leq j \leq s$ of the elements c_j of $\mathbb{F}_{q^k}^n$, for any programs T, T' , we let

$$(T, T')_i^k = \left| \{X \in \mathbb{F}_{q^k}^m \mid T(X) = c_i \wedge T'(X) = \vec{0}\} \right|$$

Then, if we denote $\overrightarrow{(T, T')^k} = (T, T')_1^k, \dots, (T, T')_s^k$, that characterizes the distribution of $T|T'$, as $P_1 \mid P_2 \approx_{q^\infty} Q_1 \mid Q_2$ if and only if $\forall k \in \mathbb{N}$. $\overrightarrow{(P_1, P_2)^k} = \overrightarrow{(Q_1, Q_2)^k}$. Using the classical inner product $\vec{x} \cdot \vec{y} = \sum_i x_i y_i$, for any k and programs $U, V, U', H' \in \overline{\mathcal{P}}_q$, we have:

$$\begin{aligned} N_k((U - V\sigma, U', V')) & \\ = \left| \left\{ X, X' \in \mathbb{F}_{q^k}^m \mid \begin{array}{l} U(X) = V(X') \\ \wedge (U'(X), V'(X)) = \vec{0} \end{array} \right\} \right| & \\ = \sum_{c \in \mathbb{F}_{q^k}^n} \left| \{X \in \mathbb{F}_{q^k}^m \mid U(X) = c \wedge U'(X) = \vec{0}\} \right| & \\ \quad \times \left| \{X \in \mathbb{F}_{q^k}^m \mid V(X) = c \wedge V'(X) = \vec{0}\} \right| & \\ = \sum_i \overrightarrow{(U, U')^k}_i \cdot \overrightarrow{(V, V')^k}_i & \\ = \overrightarrow{(U, U')^k} \cdot \overrightarrow{(V, V')^k} & \end{aligned}$$

Using scalar operations, we have that:

$$\begin{aligned} N_k(Q_1 - Q_1\sigma, Q_2, Q_2\sigma) &= N_k(P_1 - Q_1\sigma, P_2, Q_2\sigma) \\ &= N_k(P_1 - P_1\sigma, P_2, P_2\sigma) \\ \Leftrightarrow \overrightarrow{(Q_1, Q_2)^k} \cdot \overrightarrow{(Q_1, Q_2)^k} &= \overrightarrow{(P_1, P_2)^k} \cdot \overrightarrow{Q^k} \\ &= \overrightarrow{(P_1, P_2)^k} \cdot \overrightarrow{(P_1, P_2)^k} \\ \Leftrightarrow \overrightarrow{((P_1, P_2)^k - \overrightarrow{Q^k})} \cdot \overrightarrow{((P_1, P_2)^k - (Q_1, Q_2)^k)} &= \vec{0} \\ \Leftrightarrow \overrightarrow{(P_1, P_2)^k} &= \overrightarrow{(Q_1, Q_2)^k} \end{aligned}$$

Hence,

$$\begin{aligned} \forall k \in \mathbb{N}. \forall c \in \mathbb{F}_{q^k}^n. & \\ \left| \{X \in \mathbb{F}_{q^k}^m \mid P_1(X) = c \wedge P_2(X) = \vec{0}\} \right| & \\ = \left| \{X \in \mathbb{F}_{q^k}^m \mid Q_1(X) = c \wedge Q_2(X) = \vec{0}\} \right| & \\ \Leftrightarrow & \\ \forall k \in \mathbb{N}. N_k((Q_1 - Q_1\sigma, Q_2, Q_2\sigma)) & \\ = N_k((P_1 - Q_1\sigma, P_2, Q_2\sigma)) & \\ = N_k((P_1 - P_1\sigma, P_2, P_2\sigma)) & \end{aligned}$$

This concludes the proof, as for all U, V ,

$$Z(U) = Z(V) \Leftrightarrow \forall k. N_k(U) = N_k(V)$$

□

In [6], we provide a variant of this result for the specific case of verifying if a program follows the uniform distribution over all extensions, where only one computation of a local zeta function is required.

Using the complexity for the computation of the local zeta function provided by [19, Corollary 2] we obtain the following corollary.

Corollary 4.6. *q^∞ -equivalence and q^∞ -conditional equivalence restricted to arithmetic programs are in EXP.*

We now wish to remove conditionals, in order to reduce equivalence for programs with conditional to arithmetic programs (which are simply tuples of polynomials). To remove the conditionals, the first idea is to use a classical encoding inside finite fields: $\llbracket \text{if } B \neq 0 \text{ then } P_1^t \text{ else } P_1^f \rrbracket^{q^k} = \llbracket P_1^f + B^{q^k-1}(P_1^t - P_1^f) \rrbracket^{q^k}$. This works nicely as B^{q^k-1} is equal to 0 if $B = 0$, else to 1. However, for the universal case, we need to have an encoding which does not depend on the size of the field, i.e., it must be independent of k . The key idea is that for any variable t and polynomial B :

$$(B(Bt - 1) = 0 \wedge t(Bt - 1) = 0) \Leftrightarrow t = B^{q^k-2}$$

We can then encode conditionals as a multiplication by some fresh variable t , where t is conditioned by the previous conditions. An induction on the number of conditionals yields our second lemma.

Lemma 4.7. *For any $k \in \mathbb{N} \cup \{\infty\}$, q^k -conditional equivalence restricted to programs without failures reduces in exponential time to q^k -conditional equivalence restricted to arithmetic programs.*

Recall that failures define the probabilistic semantics by normalization. And for instance, for some program (if $b = 0$ then P_1 else $\perp$, P_2) where P_1 and P_2 do not fail and b is a polynomial, for any $\vec{c}$, we have:

$$\begin{aligned} & \llbracket (\text{if } b = 0 \text{ then } P_1 \text{ else } \perp, P_2) \rrbracket^{q^k}(\vec{c}, \vec{0}) \\ &= \frac{\mathbb{P}\{P_1 = \vec{c} \wedge P_2 = \vec{0} \wedge b = 0\}}{\mathbb{P}\{b = 0\}} \end{aligned}$$

Handling this division by itself would be difficult if we wanted to compute the distribution. However, in our setting, we are comparing the equality of two distributions, so we can simply multiply on both side by the denominator, and try to express once again all factors as an instance of conditional equivalence. We will be able to push inside conditional equivalence some probabilities, as $\llbracket P \rrbracket^{q^k}(\vec{c}) \times \mathbb{P}\{b = 0\} = \llbracket P, b \rrbracket^{q^k}(\vec{c}, 0)$ when all variables in b do not appear in P .

As an illustration of how to remove the failures, with some program Q , we have:

$$\begin{aligned} & \text{if } b = 0 \text{ then } P_1 \text{ else } \perp \mid P_2 \approx_{q^k} Q \mid 0 \\ & \Leftrightarrow \forall \vec{c}. \llbracket P_1, P_2, b \rrbracket^{q^k}(\vec{c}, \vec{0}) = \mathbb{P}\{b = 0\} \llbracket Q \rrbracket^{q^k}(\vec{c}) \end{aligned}$$

To reduce to an instance of conditional equivalence, the issue is that we need to express as an equality the disequality

$b \neq 0$. With some fresh variable t , multiplying by $\mathbb{P}\{\neg(b = 0)\}$ or conditioning on $tb - 1 = 0$ is equivalent, as b has an inverse if and only if it is different from zero. We can thus have:

$$\begin{aligned} & \text{if } b = 0 \text{ then } P_1 \text{ else } \perp \mid P_2 \approx_{q^k} Q \mid 0 \\ & \Leftrightarrow P_1 \mid P_2, b \approx_{q^k} Q \mid tb - 1 \end{aligned}$$

Using those techniques, we obtain:

Lemma 4.8. *For any $k \in \mathbb{N} \cup \{\infty\}$, q^k -conditional equivalence reduces to q^k -conditional equivalence restricted to programs without failures in exponential time.*

The previous Lemmas allows us to conclude.

Theorem 4.9. *q^∞ -equivalence and q^∞ -conditional equivalence are in 2-EXP.*

And using once again [6], we obtain the same complexity results for the independence problem.

Corollary 4.10. *q^∞ -conditional independence is in 2-EXP.*

Moreover, we can also extend the lower bound obtained for q -equivalence.

Lemma 4.11. *q -equivalence reduces in polynomial time to q^∞ -equivalence.*

4.3 Bounded Universal Simulatability

Simulation-based proofs [21] are one main cornerstone of cryptography. Informally, simulation-based proofs consider a real and an ideal world, and require showing the existence of a simulator, such that no adversary can distinguish the composition of the simulator and of the ideal world from the real world. This can be modelled in our context by requiring the existence of a program S (the simulator) such that “plugging in” the ideal world into S is equivalent to the real world. In this section, we consider a simpler task, where the size of the simulator is bounded. Given a program C , we denote $\deg(C)$ the maximum degree of a program, i.e the maximum degree of any polynomial appearing in C (the degree of a polynomial is the maximum over the sum of the degrees of each monomial).

Definition 4.12. Let $P, Q \in \mathcal{P}_q(I, R)$, R' such that $\#R = \#R'$ and $l \in \mathbb{N}$. We denote $P \sqsubseteq_{q^{[l]}}^l Q$, if there exists $S \in \mathcal{P}_q(\{i_1, \dots, i_n\}, R')$ such that $\deg(S) \leq l$, and $S[\mathcal{Q}/\vec{i}] \approx_{q^{[l]}} P$.

The associated decision problem is:

l, q -simulatability
INPUT: $P, Q \in \mathcal{P}_q(I, R)$.
$P \sqsubseteq_{q^{[l]}}^l Q$?

Thanks to the bound on the degree coming from l , we can easily obtain a bound on the number of such possible contexts. This is shown in [6]. From the bound on the number of contexts and the decidability of universal equivalence, one can derive the decidability of bounded simulatability.

Theorem 4.13. *l, q -simulatability is decidable.*

As a lower bound, we prove that l, q -simulatability is as hard as universal equivalence:

Lemma 4.14. *For any $l \in \mathbb{N}$, $k \in \mathbb{N} \cup \{\infty\}$, q^k -equivalence reduces in polynomial time to l, q^k -simulatability.*

We conclude this section by noting that our notion of bounded simulatability is more restricted than the general paradigm of simulation-based proofs but could be a good starting point for automating simulation-based proofs.

4.4 Universal zero-majority without inputs

For arbitrary programs, we reduce q^∞ -0-majority to the POSITIVITY problem. We recall that a Linear Recurrence Sequence (LRS) is an infinite sequence of reals $u = u_1, u_2, \dots$ such that there exist real constants $a_1, \dots, a_k$ such that for all $n \geq 0$,

$$u_{n+k} = a_1 u_{n+k-1} + \dots + a_k u_n$$

The order of a LRS $u = u_1, u_2, \dots$ is the smallest k such that the equation above holds. A LRS $u = u_1, u_2, \dots$ is positive if $u_n \geq 0$ for every $n \in \mathbb{N}$. The positivity problem consists in deciding whether a LRS is positive.

We use the fact that from a local zeta function, which is rational, we can obtain a Linear Recurrence Sequence. Then, by considering the POSITIVITY of the LRS obtained by subtracting two local zeta function, we actually check if the coefficients of the first one are always greater than the second one. We remark that the complexity of the problem strongly relies on the presence of multiplications, as for q^∞ -equivalence. Indeed, in the linear case, majority implies equivalence and we obtain the following.

Lemma 4.15. *q^∞ -0-majority restricted to linear programs is in PTIME.*

The general case has yet to be proven decidable.

Theorem 4.16. *q^∞ -0-majority for inputless programs reduces in exponential time to POSITIVITY.*

Proof. Let $P, Q \in \overline{\mathcal{P}}_q(\emptyset, R)$. We assume without loss of generality that we only have to consider arithmetic programs, using the same simplifications for observe and conditionals as we did for universal equivalence.

Recall that for any P , the local zeta function $Z(P)$ (over indeterminate T) is rational thanks to Weil's conjecture [12]. Moreover, with $N_k(P) = \left| \{X \in \mathbb{F}_{q^k}^m \mid P(X) = 0\} \right|$, we have that:

$$\frac{d}{dT} \log(Z(P)) = \frac{Z'(P)}{Z(P)} = \sum_k N_k(P) T^k$$

Let us call $\tilde{Z}(P) = \sum_k N_k(P) T^k$, which is also a rational function as Z is (and so is Z'). As the coefficients of $\tilde{Z}(P) - \tilde{Z}(Q)$ are $N_k(P) - N_k(Q)$, we have that $\forall k, N_k(P) \geq N_k(Q)$ if and only if $\tilde{Z}(P) - \tilde{Z}(Q)$ only has positive coefficients. It

is well known that the coefficients of the Taylor serie of a rational function form a LRS (see e.g. [13]). This means that the coefficients of $\tilde{Z}(P) - \tilde{Z}(Q)$ form an LRS, which we denote z^{PQ} . We finally get that $Q <_{q^\infty} P$ if and only if $\forall n, z_n^{PQ} \geq 0$. $\square$

This reduction can also be applied with the generalization of [18], and thus, for any two events about programs over finite fields, one can, given an oracle for the POSITIVITY problem, decide if the probability of the first event is greater than the second one for all extensions of the base field.

Similarly to the equivalence case, we can derive some hardness from the non universal case, but we do not obtain any completeness result.

Lemma 4.17. *2^∞ -0-majority is PP-hard.*

5 Conclusion

We have introduced universal equivalence and majority problems and studied their complexity and decidability. Our work could notably be used as a building block to design a decidable logic for universal probabilistic program verification. It leaves several questions of interest open:

- the exact complexity of universal equivalence is open. It is even unknown whether the universal problem is strictly harder than the non-universal one;
- the decidability of universal majority is open. The decidability of POSITIVITY would yield decidability of universal 0-majority and equivalently, undecidability of universal majority would also solve negatively the POSITIVITY problem;
- the decidability of universal approximate equivalence is open. Approximate equivalence asks whether the statistical distance between the distributions of two programs is negligible in k . This notion has direct applications in provable security.

Acknowledgements. We wish to thank the anonymous reviewers for their useful comments, as well as Martin Grohe, Bruce Kapron, David Mestel, Joël Ouaknine, Pierre-Jean Spaenlehauer, Emmanuel Thomé and Mehdi Tibouchi for interesting discussions. We are grateful for the support by the ERC under the EU's Horizon 2020 research and innovation program (grant agreement No 645865-SPOOC) as well as ONR (grant N000141512750).

References

- [1] James Ax. 1968. The elementary theory of finite fields. *Annals of Mathematics* 88, 2 (1968), 239–271.
- [2] Gilles Barthe, Rohit Chadha, Vishal Jagannath, A. Prasad Sistla, and Mahesh Viswanathan. 2019. Automated Methods for Checking Differential Privacy. *CoRR* abs/1910.04137 (2019). [arXiv:1910.04137](https://arxiv.org/abs/1910.04137) [http://arxiv.org/abs/1910.04137](https://arxiv.org/abs/1910.04137)
- [3] Gilles Barthe, Marion Daubignard, Bruce Kapron, Yassine Lakhnech, and Vincent Laporte. 2010. On the equality of probabilistic terms. In

- 16th International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR'10) (Lecture Notes in Computer Science), Vol. 6355. Springer, 46–63.
- [4] Gilles Barthe, Benjamin Grégoire, Charlie Jacomme, Steve Kremer, and Pierre-Yves Strub. 2019. Symbolic Methods in Computational Cryptography Proofs. In *32nd IEEE Computer Security Foundations Symposium (CSF'19)*. IEEE Computer Society, 136–151.
- [5] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. 2009. Formal certification of code-based cryptographic proofs. *ACM SIGPLAN Notices* 44, 1 (2009), 90–101.
- [6] Gilles Barthe, Charlie Jacomme, and Steve Kremer. 2020. Universal equivalence and majority on probabilistic programs over finite fields. (2020). <https://hal.inria.fr/hal-02552287>
- [7] Benjamin Bichsel, Timon Gehr, and Martin Vechev. 2018. Fine-grained Semantics for Probabilistic Programs. In *27th European Symposium on Programming (ESOP'18) (Lecture Notes in Computer Science)*, Vol. 10801. Springer, 145–185.
- [8] Enrico Bombieri. 1966. On exponential sums in finite fields. *American Journal of Mathematics* 88, 1 (1966), 71–105.
- [9] Brent Carmer and Mike Rosulek. 2016. Lincrypt: a model for practical cryptography. In *36th Annual International Cryptology Conference (CRYPTO'16) (Lecture Notes in Computer Science)*, Vol. 9816. Springer, 416–445.
- [10] Dmitry Chistikov, Andrzej S Murawski, and David Purser. 2019. Asymmetric Distances for Approximate Differential Privacy. In *30th International Conference on Concurrency Theory (CONCUR 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [11] Benny Chor, Oded Goldreich, Johan Håstad, Joel Friedman, Steven Rudich, and Roman Smolensky. 1985. The Bit Extraction Problem of t -Resilient Functions (Preliminary Version). In *26th Annual Symposium on Foundations of Computer Science (FOCS'85)*. IEEE Computer Society, 396–407.
- [12] Bernard Dwork. 1960. On the rationality of the zeta function of an algebraic variety. *American Journal of Mathematics* 82, 3 (1960), 631–648.
- [13] Graham Everest, Alf van der Poorten, Igor Shparlinski, and Thomas Ward. 2002. Exponential functions, linear recurrence sequences, and their applications.
- [14] Matthew Fredrikson and Somesh Jha. 2014. Satisfiability modulo counting: A new approach for analyzing privacy properties. In *Joint Meeting of the 23rd Annual Conference on Computer Science Logic (CSL) and the 29th ACM/IEEE Symposium on Logic in Computer Science (LICS)*. ACM, 1–10.
- [15] Marco Gaboardi, Kobbi Nissim, and David Purser. 2019. The Complexity of Verifying Circuits as Differentially Private. *CoRR* abs/1911.03272 (2019). arXiv:1911.03272 <http://arxiv.org/abs/1911.03272>
- [16] William Andrew Johnson. 2016. *Fun with fields*. Ph.D. Dissertation. UC Berkeley.
- [17] Charanjit S Jutla and Arnab Roy. 2012. Decision procedures for simulatability. In *17th European Symposium on Research in Computer Security (ESORICS'12) (Lecture Notes in Computer Science)*, Vol. 7459. Springer, 573–590.
- [18] Catarina Kiefe. 1976. Sets definable over finite fields: their zeta-functions. *Trans. Amer. Math. Soc.* 223 (1976), 45–59.
- [19] Alan GB Lauder and Daqing Wan. 2006. Counting points on varieties over finite fields of small characteristic. *arXiv preprint math/0612147* (2006).
- [20] Axel Legay, Andrzej S Murawski, Joël Ouaknine, and James Worrell. 2008. On automated verification of probabilistic programs. In *14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'08) (Lecture Notes in Computer Science)*, Vol. 4963. Springer, 173–187.
- [21] Yehuda Lindell. 2017. How to Simulate It - A Tutorial on the Simulation Proof Technique. In *Tutorials on the Foundations of Cryptography*. Yehuda Lindell (Ed.). Springer International Publishing, 277–346. https://doi.org/10.1007/978-3-319-57048-8_6
- [22] Michael L Littman, Judy Goldsmith, and Martin Mundhenk. 1998. The computational complexity of probabilistic planning. *Journal of Artificial Intelligence Research* 9 (1998), 1–36.
- [23] Stefan Maubach. 2001. The automorphism group over finite fields. (2001).
- [24] Scott McCallum and Volker Weispfenning. 2012. Deciding polynomial-transcendental problems. *Journal of Symbolic Computation* 47, 1 (2012), 16–31.
- [25] Maurice Mignotte, Tarlok Nath Shorey, and Robert Tijdeman. 1984. The distance between terms of an algebraic recurrence sequence. *Journal für die reine und angewandte Mathematik* 349 (1984), 63–76.
- [26] Andrzej S. Murawski and Joël Ouaknine. 2005. On Probabilistic Program Equivalence and Refinement. In *CONCUR 2005 - Concurrency Theory, 16th International Conference, CONCUR 2005, San Francisco, CA, USA, August 23-26, 2005, Proceedings (Lecture Notes in Computer Science)*, Martín Abadi and Luca de Alfaro (Eds.), Vol. 3653. Springer, 156–170. https://doi.org/10.1007/11539452_15
- [27] Tobias Nipkow. 1990. Unification in Primal Algebras, Their Powers and Their Varieties. *J. ACM* 37, 4 (Oct. 1990), 742–776.
- [28] Joël Ouaknine and James Worrell. 2012. Decision problems for linear recurrence sequences. In *6th International Workshop on Reachability Problems (RP'12) (Lecture Notes in Computer Science)*, Vol. 7550. Springer, 21–28.
- [29] Joël Ouaknine and James Worrell. 2014. Positivity problems for low-order linear recurrence sequences. In *25th ACM-SIAM Symposium on Discrete Algorithms (SODA'14)*. Society for Industrial and Applied Mathematics, 366–379.
- [30] Jacobo Torán. 1988. An oracle characterization of the counting hierarchy. In *3rd Annual Structure in Complexity Theory Conference*. 213–223.
- [31] Jacobo Torán. 1991. Complexity classes defined by counting quantifiers. *J. ACM* 38 (1991), 753–774.
- [32] NK Vereshchagin. 1985. The problem of appearance of a zero in a linear recurrence sequence. *Mat. Zametki* 38, 2 (1985), 609–615.
- [33] André Weil. 1949. Numbers of solutions of equations in finite fields. *Bulletin of the AMS* (1949).