Déjà Q: Using Dual Systems to Revisit q-Type Assumptions

Melissa Chase[1] and Sarah Meiklejohn[2⋆]

[1] Microsoft Research Redmond
melissac@microsoft.com
[2] UC San Diego
smeiklej@cs.ucsd.edu

Abstract. After more than a decade of usage, bilinear groups have established their place in the cryptographic canon by enabling the construction of many advanced cryptographic primitives. Unfortunately, this explosion in functionality has been accompanied by an analogous growth in the complexity of the assumptions used to prove security. Many of these assumptions have been gathered under the umbrella of the "uber-assumption," yet certain classes of these assumptions — namely, *q-type* assumptions — are stronger and require larger parameter sizes than their static counterparts. In this paper, we show that in certain bilinear groups, many classes of q-type assumptions are in fact implied by subgroup hiding (a well-established, static assumption). Our main tool in this endeavor is the *dual-system* technique, as introduced by Waters in 2009. As a case study, we first show that in composite-order groups, we can prove the security of the Dodis-Yampolskiy PRF based solely on subgroup hiding and allow for a domain of arbitrary size (the original proof only allowed a logarithmically-sized domain). We then turn our attention to classes of q-type assumptions and show that they are implied — when instantiated in appropriate groups — solely by subgroup hiding. These classes are quite general and include assumptions such as q-SDH. Concretely, our result implies that every construction relying on such assumptions for security (e.g., Boneh-Boyen signatures) can, when instantiated in appropriate composite-order bilinear groups, be proved secure under subgroup hiding instead.

1 Introduction

For the past decade, bilinear groups — i.e., groups equipped with a bilinear map, or pairing — have allowed for the efficient construction of a wide variety of advanced cryptographic primitives, including (but by no means limited to): signatures [**?**,**?**,**?**,**?**], group signatures [**?**,**?**,**?**], zero-knowledge proofs [**?**,**?**], (hierarchical) identity-based encryption [**?**,**?**,**?**,**?**], and functional and attribute-based encryption [**?**,**?**,**?**]. As such, pairings are now used as a standard general-purpose tool in cryptographic constructions.

⋆ Work done as an intern at Microsoft Research Redmond

Unfortunately, this growth in the complexity of cryptographic primitives has been accompanied by an analogous growth in the complexity of the assumptions required to prove security. While assumptions such as Bilinear Diffie Hellman (BDH) [**?**] and Decision Linear [**?**] have become relatively standard, the use of pairings has also ushered in various classes of assumptions such as *q-type assumptions*, in which the size of the assumption grows dynamically, or *interactive assumptions*, in which the adversary is given access to some oracle(s). For example, in the q-DBDHI (Decisional Bilinear Diffie Hellman Inversion) assumption, the adversary is given $(g, g^x, g^{x^2}, \ldots, g^{x^q})$ and is asked to produce $e(g,g)^{1/x}$. While the "uber-assumption" [**?**,**?**] generalizes many q-type assumptions (as well as many static assumptions) and provides a lower bound for their security in the generic group model [**?**], such assumptions nevertheless remain less understood than their static counterparts.

Beyond the understanding of such assumptions, the fact that they scale asymptotically with the security of the scheme can be problematic. In a reduction, the value of q is frequently tied to the number of queries that the adversary makes to an oracle. As a result, q must scale with some parameter of the system; e.g., for identity-based encryption, q must be at least as big as the number of parties that the adversary is able to corrupt. As it is typically the case that an assumption parameterized by q' implies the same assumption parameterized by q for $q' > q$ (as the assumption parameterized by q' gives out strictly more information), this means that the assumption gets stronger as the adversary is able to corrupt more parties. In some cases, this correlation is more striking. For example, Dodis and Yampolskiy [**?**] use the $2^{a(\lambda)}$-DBDHI assumption to prove the security of their pseudorandom function (PRF), where $a(\lambda)$ is the size of the domain of the PRF (and λ is the security parameter); as a result, the domain is restricted to be of logarithmic size. This correlation is furthermore not always an artifact of proof techniques, as Jao and Yoshida [**?**] showed that Boneh-Boyen signatures were in fact equivalent to the q-SDH assumption that they rely on for security. Finally, Cheon [**?**] showed that the time required to recover a secret key scales inversely with the size of q, so that if recovering a secret key takes time t when using $q = 1$ (e.g, it takes t steps to recover x given g and g^x), then it takes time $t/\sqrt{q}$ in the general case (e.g., given $(g, g^x, \ldots, g^{x^q})$). This means that constructions rely on asymptotically stronger assumptions to obtain stronger security guarantees, so the parameters must grow appropriately in order to maintain a constant level of security (e.g., 128-bit security).

On the positive side, one technique that has proved particularly effective at avoiding q-type assumptions — and boosting security as a result — is the *dual-system* technique, which was introduced by Waters [**?**] in 2009 and has been used extensively since [**?**,**?**,**?**,**?**,**?**,**?**]. Briefly, this technique takes advantage of *subgroup hiding* in bilinear groups [**?**]; i.e., the assumption, in a group of composite order $N = p_1 p_2$, that a random element of the full group is indistinguishable from a random element of order p_1. (Subgroup hiding can also be defined, albeit in a more complex way, for vector spaces over prime-order bilinear groups.) Using this core assumption, the dual-system technique begins with a scheme in a particular

subgroup (for concreteness, the subgroup of elements of order p_1); i.e., a scheme in which all elements are contained solely within the subgroup. To prove security, a "shadow" copy of the original scheme is first added in a new subgroup (e.g., the subgroup of order p_2); the addition of this shadow copy goes unnoticed by subgroup hiding. Using a property called *parameter hiding* [**?**], this shadow copy is then randomized, so the value in the additional subgroup is now unstructured; in Waters' terminology, this object is now *semi-functional*. This randomness is then pushed back into the original subgroup, again using subgroup hiding, and is used to blind the structure of the original scheme; e.g., in an IND-CPA game it can be used to obscure all information about the challenge message.

Our contributions. In this paper, we expand the usage of the dual-system technique. Rather than work at the level of constructions, we show directly that many q-type assumptions can be implied — with a crucial looseness of q — by subgroup hiding. In some sense, we thus interpret previous usages as *absorbing* rather than avoiding q-type assumptions, and believe our work takes a (perhaps surprising) step in expanding the power of the dual-system technique.

As a first exercise, we prove in Section 3 that the Dodis-Yampolskiy PRF — unmodified, but instantiated in a composite-order group — can be proved secure using only the subgroup hiding assumption. Because of the limitations (described above) in the original security proof, our result not only uses a static assumption, but also boosts security to allow for domains of arbitrary size, which is useful in and of itself for the many applications of the Dodis-Yampolskiy PRF [**?**,**?**,**?**,**?**].

Next, in Section 4, we look beyond cryptographic primitives and instead focus directly on the underlying assumptions, and in particular on the class of q-type assumptions that are instantiations of the uber-assumption. Here we show that many instantiations of the uber-assumption can be reduced — following a modified version of the dual-system technique, which still assumes subgroup hiding — to instantiations that are significantly weaker; in fact, in many cases we can reduce to an assumption so weak that it actually holds by a statistical argument. As examples, we revisit a number of well-known q-type assumptions. By applying our general theorem to these assumptions, we can reduce them to assumptions in which all secret information (e.g., the exponent x in q-DBDHI) is statistically hidden, so an adversary can do no better than a random guess and the security of the entire assumption collapses down to subgroup hiding.

Finally, in Section 5, we discuss the concrete implications of our work; i.e., in which concrete bilinear settings the abstract requirements of the dual-system technique (namely, subgroup hiding and parameter hiding) can be expected to hold. Due to current limitations in the parameter hiding supported by prime-order bilinear groups, our results can most generally be applied in asymmetric composite-order bilinear groups [**?**,**?**].

Putting it all together, we obtain the following concrete results:

– In a composite-order group (such as the target group of a composite-order pairing, or any composite-order elliptic curve group without a pairing), sub-

group hiding implies any q-type assumption where the exponents are linearly independent rational functions.
- In an asymmetric composite-order bilinear group, subgroup hiding implies any q-type assumption where the exponents are linearly independent rational functions and the adversary must compute a value in the source group.

Related work. As mentioned above, the dual-system technique was first introduced by Waters in 2009 [**?**], and was applied subsequently to achieve a wide variety of results [**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**,**?**], all involving randomized public-key primitives (e.g., identity-based encryption) in bilinear groups.

To the best of our knowledge, we are the first to systematically apply the dual-system technique directly to assumptions, and in particular to q-type assumptions. Boneh, Boyen, and Goh [**?**] analyzed the security of the uber-assumption — which includes many q-type assumptions — in the generic group model, and derived generic lower bounds on the runtime of an adversary that could break the uber-assumption; this work was later extended by Jager and Rupp [**?**], who showed the equivalence of many assumptions in the *semi-generic* group model. Our result is somewhat orthogonal to theirs, as we seek to show that in certain concrete (i.e., non-generic) settings these assumptions actually reduce to subgroup hiding. Anecdotally, several results use the dual-system technique to eliminate the requirement on q-type assumptions for specific primitives or constructions: Gerbush et al. [**?**] obtained Camenisch-Lysyanskaya signatures under static assumptions, as opposed to the interactive LRSW assumption; Attrapadung and Libert achieved the first identity-based broadcast encryption scheme with short ciphertexts [**?**]; and the original result of Waters [**?**] achieved the first secure HIBE under non-q-type assumptions.

2 Definitions and Notation

2.1 Preliminaries

If x is a binary string then $|x|$ denotes its bit length. If S is a finite set then $|S|$ denotes its size and $x \xleftarrow{\$} S$ denotes sampling a member uniformly from S and assigning it to x. $\lambda \in \mathbb{N}$ denotes the security parameter and 1^λ denotes its unary representation.

Algorithms are randomized unless explicitly noted otherwise. "PT" stands for "polynomial-time." By $y \leftarrow A(x_1, \ldots, x_n; R)$ we denote running algorithm A on inputs $x_1, \ldots, x_n$ and random coins R and assigning its output to y. By $y \xleftarrow{\$} A(x_1, \ldots, x_n)$ we denote $y \leftarrow A(x_1, \ldots, x_n; R)$ for some random coins R. By $[A(x_1, \ldots, x_n)]$ we denote the set of values that have positive probability of being output by A on inputs $x_1, \ldots, x_n$. Adversaries are algorithms.

We use games in definitions of security and in proofs. A game G has a MAIN procedure whose output is the output of the game. $\Pr[\mathsf{G}]$ denotes the probability that this output is true.

2.2 Bilinear groups

We refer to a *bilinear group* as a tuple $\mathbb{G} = (N, G, H, G_T, e)$, where N can be either prime or composite, $|G| = |H| = kN$ and $|G_T| = \ell N$ for some $k, \ell \in \mathbb{N}$, and $e : G \times H \to G_T$ is a bilinear map, meaning it is (1) efficiently computable; (2) satisfies bilinearity: $e(x^a, y^b) = e(x, y)^{ab}$ for all $x \in G$, $y \in H$, and $a, b \in \mathbb{Z}/N\mathbb{Z}$; and (3) satisfies non-degeneracy: if $e(x, y) = 1$ for all $y \in H$ then $x = 1$ and if $e(x, y) = 1$ for all $x \in G$ then $y = 1$. When G and H are cyclic, we may include in $\mathbb{G}$ generators g and h of G and H respectively, and when the groups G and H decompose into subgroups $G = G_1 \oplus G_2$ and $H = H_1 \oplus H_2$, we may additionally include descriptions of these subgroups and/or their generators. In what follows, we use $\mathsf{BilinearGen}$ to denote the algorithm by which bilinear groups are generated, and provide it with an argument n that specifies the number of subgroups.

There are two additional structural properties of bilinear groups that are exploited in the dual-system technique: subgroup hiding and parameter hiding. Subgroup hiding is a computational assumption that requires that, if G (respectively H) decomposes into two subgroups, then distinguishing between a random element of the full group and a random element of one of the subgroups should be hard. (This is actually the specific simple case of subgroup hiding originally introduced by Boneh, Goh, and Nissim [**?**]; more general definitions exist as well [**?**,**?**].)

Assumption 2.1 (Subgroup hiding) *For a bilinear group* $\mathbb{G} = (N, G, H, G_T, e, g_1, g_2, h_1, h_2)$, subgroup hiding *holds in* $\mathbb{G}$ *if no PT adversary* $\mathcal{A}$ *has a non-negligible chance of distinguishing a random element of the subgroup* G_1 *from a random element of the group* G*; formally, define* $\mathbf{Adv}^{sgh}_{\mathcal{A}}(\lambda) = 2Pr[SGH^{\mathcal{A}}_{\mu}(\lambda)] - 1$*, where* $SGH^{\mathcal{A}}_{\mu}(\lambda)$ *is defined as follows for* $\mu \subseteq \{g_1, g_2, h_1, h_2\}$*:*

$\underline{\text{MAIN } SGH^{\mathcal{A}}_{\mu}(\lambda)}$
$b \xleftarrow{\$} \{0,1\}$; $(N, G, H, G_T, e, g_1, g_2, h_1, h_2) \xleftarrow{\$} \mathsf{BilinearGen}(1^\lambda, 2)$
if $(b = 0)$ *then* $T \xleftarrow{\$} G$
if $(b = 1)$ *then* $T \xleftarrow{\$} G_1$
$b' \xleftarrow{\$} \mathcal{A}((N, G, H, G_T, e), \mu, T)$
return $(b' = b)$

Then subgroup hiding *holds with respect to the auxiliary information* μ *if for all PT adversaries* $\mathcal{A}$ *there exists a negligible function* $\nu(\cdot)$ *such that* $\mathbf{Adv}^{sgh}_{\mathcal{A}}(\lambda) < \nu(\lambda)$.

There are often limits to the auxiliary information that can be provided to $\mathcal{A}$; e.g., if $\mathcal{A}$ is attempting to distinguish $T = g_1^r$ from $T = g^r$ for $r \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$ and has access to a canceling pairing $e(\cdot, \cdot)$ — i.e., a pairing such that $e(G_1, H_2) = e(G_2, H_1) = 1$ — and $h_2 \in \mu$, it can easily distinguish between these elements by checking if $e(T, h_2) = 1$ or not. Thus, if an adversary is trying to distinguish

between a random element of G_1 and a random element of $G_1 \oplus G_2$ (analogously, if it is trying to distinguish between G_2 and $G_1 \oplus G_2$), the problem becomes easy if μ includes h_2 (analogously, h_1).

Parameter hiding, unlike subgroup hiding, is a statistical property of the group that allows certain distributions across subgroups to be independent. In composite-order groups, for example, the Chinese Remainder Theorem tells us that the values of $x \bmod p_1$ and $x \bmod p_2$ are independent, so that given g_1^x, the value of g_2^x is unconstrained. In prime-order groups, Lewko [**?**] demonstrated how to support parameter hiding with respect to linear functions; i.e., how — using appropriate constructions of G_1 and G_2 — the distribution of g_2^{ax} and g_2^r for $a, r \xleftarrow{\$} \mathbb{F}_p$ is identical, even given x and g_1^a. The first formal notion of parameter hiding with respect to these linear functions was later given by Lewko and Meiklejohn [**?**]; we generalize their notion as follows:

Definition 2.1 (Parameter hiding). *For a bilinear group $\mathbb{G} = (N, G, H, G_T, e)$,* parameter hiding *holds in G with respect to a family of functions $\mathcal{F}$ if the distribution $\{g_1^{f(x_1,\ldots,x_n)} g_2^{f(x_1,\ldots,x_n)}\}_{f\in\mathcal{F}}$ is identical to $\{g_1^{f(x_1,\ldots,x_n)} g_2^{f(x'_1,\ldots,x'_n)}\}_{f\in\mathcal{F}}$ for $g_1 \xleftarrow{\$} G_1$, $g_2 \xleftarrow{\$} G_2$, and $x_1, x'_1, \ldots, x_n, x'_n \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$. (And holds analogously in H using h_1 and h_2.)*

As a very simple example, if $\mathcal{F} = \{1, x_1\}$, for $g_1 \xleftarrow{\$} G_1$, $g_2 \xleftarrow{\$} G_2$, and $x_1, x'_1 \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$, the distributions $(g_1 g_2, g_1^{x_1} g_2^{x_1})$ and $(g_1 g_2, g_1^{x_1} g_2^{x'_1})$ are identical.

We also define a somewhat weaker condition, which requires distributions that are statistically close for any (potentially adaptively chosen) polynomial-sized subset of $\mathbb{F}$.

Definition 2.2 (Adaptive parameter hiding). *For a bilinear group $\mathbb{G} = (N, G, H, G_T, e)$,* adaptive parameter hiding *holds with respect to a family of functions $\mathcal{F}$ if for all $\lambda \in \mathbb{N}$ and all adaptively chosen sets $S \subseteq \mathcal{F}$ of size poly(λ), the distribution $\{g_1^{f(x_1,\ldots,x_n)} g_2^{f(x_1,\ldots,x_n)}\}_{f\in S}$ is statistically close to $\{g_1^{f(x_1,\ldots,x_n)} g_2^{f(x'_1,\ldots,x'_n)}\}_{f\in S}$ for $g_1 \xleftarrow{\$} G_1$, $g_2 \xleftarrow{\$} G_2$, and $x_1, x'_1 \ldots, x_n, x'_n \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$.*

We use these definitions in Sections 4, and discuss the different families of functions that can be supported in different types of bilinear groups in Section 5.

2.3 Pseudorandom functions

A pseudorandom function family [**?**] F specifies the algorithms $\mathsf{F.Pg}$, $\mathsf{F.Keys}$, $\mathsf{F.Dom}$, $\mathsf{F.Rng}$, and $\mathsf{F.Ev}$. Via $fp \xleftarrow{\$} \mathsf{F.Pg}(1^\lambda)$ one generates a description fp of a function $\mathsf{F.Ev}(1^\lambda, fp)\colon \mathsf{F.Keys}(1^\lambda, fp) \times \mathsf{F.Dom}(1^\lambda, fp) \to \mathsf{F.Rng}(1^\lambda, fp)$. The evaluation algorithm $\mathsf{F.Ev}$ is PT and deterministic.

Definition 2.3. *For a function family* F *and an adversary* $\mathcal{A}$*, let* $\mathbf{Adv}^{prf}_{\mathsf{F},\mathcal{A}}(\lambda) = 2\Pr[PRF^{\mathcal{A}}_{\mathsf{F}}(\lambda)] - 1$*, where* $PRF^{\mathcal{A}}_{\mathsf{F}}(\lambda)$ *is defined as follows:*

MAIN $PRF^{\mathcal{A}}_{\mathsf{F}}(\lambda)$	*Procedure* $\mathrm{FN}_{sk}(x)$
$b \xleftarrow{\$} \{0,1\}; fp \xleftarrow{\$} \mathsf{F.Pg}(1^\lambda); sk \xleftarrow{\$} \mathsf{F.Keys}(1^\lambda, fp)$	*if* $b = 0$ $y \xleftarrow{\$} \mathsf{F.Rng}(1^\lambda, fp)$
$b' \xleftarrow{\$} \mathcal{A}^{\mathrm{FN}}(1^\lambda, fp)$	*if* $b = 1$ $y \leftarrow \mathsf{F.Ev}(1^\lambda, fp, sk, x)$
return $(b' = b)$	*return* y

Then F *is* pseudorandom *if for all PT algorithms* $\mathcal{A}$ *there exists a negligible function* $\nu(\cdot)$ *such that* $\mathbf{Adv}^{prf}_{\mathsf{F},\mathcal{A}}(\lambda) \leq \nu(\lambda)$*.*

3 Pseudorandom Functions

In this section, we explore the security of the Dodis-Yampolskiy PRF [**?**]. First, we recall the Dodis-Yampolskiy PRF, instantiated for our purposes in a group of composite order $N = p_1p_2$:[3]

- $\mathsf{F.Pg}(1^\lambda)$: Output $(N, G, H, G_T, e, g, h) \xleftarrow{\$} \mathsf{BilinearGen}(1^\lambda, 2)$. Then $\mathsf{F.Keys} = \mathsf{F.Dom} = \mathbb{Z}/N\mathbb{Z}$, and $\mathsf{F.Rng} = G_T$.
- $\mathsf{F.Ev}(1^\lambda, fp, sk, x)$: Output $y := e(g,h)^{\frac{1}{sk+x}}$. If $(sk + x)^{-1}$ is undefined in $\mathbb{Z}/N\mathbb{Z}$, then output $y := 1$.

Dodis and Yampolskiy originally showed that this was a *verifiable* random function — a more powerful primitive than a PRF, as it comes with the additional ability to prove that the PRF value was computed correctly — under the q-DBDHI assumption, which states that when given $(g, g^x, \ldots, g^{x^q})$, it should be hard to distinguish $e(g,g)^{1/x}$ from random. Their reduction, however, is quite loose: if the size of the PRF domain is $a(\lambda)$, then they use the $2^{a(\lambda)}$-DBDHI assumption, and show that $\mathbf{Adv}^{\text{pr-vrf}}_{\mathsf{F},\mathcal{A}}(\lambda) \leq 2^{a(\lambda)} \cdot \mathbf{Adv}^{2^{a(\lambda)}\text{-DBDHI}}_{\mathcal{A}}(\lambda)$, which means that the scheme is provably secure only if the domain is restricted to be of logarithmic size (i.e., its size is logarithmic in the size of the security parameter).

We instead make two minor modifications to the PRF and show that

$$\mathbf{Adv}^{\text{prf}}_{\mathsf{F},\mathcal{A}}(\lambda) \leq q \cdot \mathbf{Adv}^{\text{sgh}}_{\mathcal{A}}(\lambda)$$

for an adversary $\mathcal{A}$ that makes q queries to the PRF oracle; while the reduction is still not tight, our approach nevertheless allows for a domain of arbitrary size. Our first modification is to move the scheme into a subgroup: rather than use $e(g,h)$ for the full group generators, we switch to using $e(g_1, h_1)$, where g_1 and h_1 generate G_1 and H_1 respectively. (In a cyclic group, such as a composite-order

[3] To mirror the exposition of the original PRF, we use the target group G_T here, but note that in fact our analysis would work for *any* composite-order group in which subgroup hiding holds and there is no pairing.

target group, this could instead be accomplished using an additional application of subgroup hiding, and we can show that the function is a PRF even in the full group.) Our second modification is to use, rather than the "canonical" generator $e(g_1, h_1)$, a random generator $w_1 \in G_{T,1}$. We stress that these modifications are purely syntactical and do not fundamentally alter the spirit of the construction (and, in particular, do not affect its usage in applications). They do, however, allow us to prove the following theorem:

Theorem 3.1. *For all $\lambda \in \mathbb{N}$ and $fp \in [\mathsf{F.Pg}(1^\lambda)]$, if subgroup hiding holds in $\mathbb{G}$ and adaptive parameter hiding holds with respect to $\{\mathsf{F.Ev}(1^\lambda, fp, \cdot, x)\}_{x \in \mathsf{F.Dom}}$, then F is a pseudorandom function family.*

A proof of Theorem 3.1 can be found in the full version of the paper. Intuitively, our approach amplifies the only unknown value present in the PRF — namely, the *sk* value — as follows: first, this secret value is replicated in the $G_{T,2}$ subgroup, which is indistinguishable from the original by subgroup hiding. The secret value in the $G_{T,2}$ subgroup is then decoupled from the secret value in the $G_{T,1}$ subgroup, which is indistinguishable (in fact identical) by parameter hiding. Finally, the new secret value from the $G_{T,2}$ subgroup is moved back into $G_{T,1}$, which is again indistinguishable by subgroup hiding. At this point, we now have one additional secret value in the PRF values we return. By repeating the process, we can embed polynomially many secret values (in particular, we embed as many values as there are oracle queries), at which point we have enough entropy to argue that the values returned by the PRF are statistically indistinguishable from truly random values.

One interesting feature of our approach is that — because we are using a deterministic primitive — we do not need to follow the traditional dual-system structure and adhere to a "query hybrid," in which each query to the oracle must be treated separately. Nevertheless, we do need to add enough additional degrees of randomness to cover all of the adversary's queries, so we still end up with a looseness of q in our reduction (but where q is the number of queries, not the size of the PRF domain).

4 Reducing *q*-Type Assumptions to Subgroup Hiding

Our main result in this section is to show that — if subgroup hiding holds and parameter hiding holds with respect to certain functions in the exponent — certain q-type assumptions are equivalent to significantly weaker assumptions. In fact, these equivalent assumptions are often so weak that they hold by a purely statistical argument, so the original assumption is fully implied by subgroup hiding.

We begin by recalling the uber-assumption, which serves as an umbrella for many q-type assumptions. We then describe two approaches: roughly, the first reduces any uber-assumption to subgroup hiding, but only if the assumption gives out meaningful functions on one side of the pairing (or in the target group), and the second reduces any computational uber-assumption in the source group

to subgroup hiding. Both of our reductions incur a looseness of q in the reduction, so we can think of them as "absorbing" the factor of q from the assumption rather than eliminating it outright.

4.1 The uber-assumption

We are able to examine many q-type assumptions at the same time using the "uber-assumption" [**?**,**?**], which was first introduced by Boneh, Boyen, and Goh as a way to reason generally about a wide variety of pairing-based assumptions. They prove that if the parameters of the uber-assumption meet certain independence requirements then the assumption is hard in the generic group model, which eliminates the need to prove generic lower bounds for every individual instantiation of the assumption that is introduced. Our motivation, on the other hand, is to prove that many common instantiations of the assumption are in fact implied — assuming subgroup hiding holds in the bilinear group — by weaker versions of the assumption.

Formally, for a bilinear group $\mathbb{G} = (N, G, H, G_T, e, g, h)$ (where N can be either prime or composite) the uber-assumption is parameterized by five values: an integer $c \in \mathbb{N}$, three sets R, S, and T of polynomials over $\mathbb{Z}/N\mathbb{Z}$ (which represent the values we are given in G, H, and G_T respectively), and a polynomial f over $\mathbb{Z}/N\mathbb{Z}$. For the sets of polynomials, we write $R = \langle \rho_1(x_1, \ldots, x_c), \ldots, \rho_r(x_1, \ldots, x_c)\rangle$ and as shorthand use $\rho_i(\vec{x}) = \rho_i(x_1, \ldots, x_c)$ and $g^{R(x_1,\ldots,x_c)} = \{g^{\rho_i(\vec{x})}\}_{i=1}^{r}$ (and similarly for S and T).

Assumption 4.1 (Computational) *For an adversary $\mathcal{A}$, define* $\mathbf{Adv}^{uber}_{\mathcal{A}}(\lambda) = Pr[c\text{-}UBER^{\mathcal{A}}_{c,R,S,T,f}(\lambda)]$*, where* $c\text{-}UBER^{\mathcal{A}}_{c,R,S,T,f}(\lambda)$ *is defined as follows:*

$\underline{\text{MAIN } c\text{-}UBER^{\mathcal{A}}_{c,R,S,T,f}(\lambda)}$
$(N, G, H, G_T, e, g, h) \xleftarrow{\$} \mathsf{BilinearGen}(1^\lambda, 2)$; $x_1, \ldots, x_c \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$
$y \xleftarrow{\$} \mathcal{A}(1^\lambda, (N, G, H, G_T, e, g, h), g^{R(x_1,\ldots,x_c)}, h^{S(x_1,\ldots,x_c)}, e(g,h)^{T(x_1,\ldots,x_c)})$
return $(y = e(g,h)^{f(x_1,\ldots,x_c)})$

Then the uber-assumption holds if for all PT algorithms $\mathcal{A}$ there exists a negligible function $\nu(\cdot)$ such that $\mathbf{Adv}^{uber}_{\mathcal{A}}(\lambda) < \nu(\lambda)$.

As an example, CDH in a symmetric group G uses $c = 2$, $R = S = \langle 1, x_1, x_2\rangle$, $T = \langle 1\rangle$, and $f(x_1, x_2) = x_1 x_2$, so that given (g, g^{x_1}, g^{x_2}), it should be hard to compute $g^{x_1 x_2}$. As long as R and S both include 1, the computational uber-assumption in the target group implies the computational uber-assumption in the source group, since given $X = g^{f(\vec{x})}$ one can always compute $e(X, h) = e(g, h)^{f(\vec{x})}$.

The game d-UBER$^{\mathcal{A}}_{c,R,S,T,f}(\lambda)$ for the decisional uber-assumption is defined analogously, except rather than compute $g^{f(x_1,\ldots,x_c)}$ at the end, the adversary has only to distinguish it from random. Unlike the computational version, the decisional uber-assumption in the source group implies the decisional uber-assumption in the target group, since one can use a decider between $e(g, h)^{f(\vec{x})}$

and R_T to decide between $g^{f(\vec{x})}$ and R by computing the pairing. Furthermore, the decisional uber-assumption (in either group) implies the computational uber-assumption, since the ability to compute the target value immediately implies the ability to distinguish it from random. The strongest version of the uber-assumption, and the one we therefore choose to aim for in the next section, is the decisional assumption in either of the source groups.

4.2 A first approach: functions on one side of the pairing

Our first approach shows that certain classes of the uber-assumption are equivalent to significantly weaker classes, and that in fact these weaker classes are so weak that the assumption holds by a statistical argument. The subclass of uber-assumptions we cover includes q-type assumptions such as exponent q-SDH (defined above), and implies that any schemes that currently rely on such assumptions can be instantiated so that they rely solely on subgroup hiding.

Our only modifications to the parameters of the uber-assumption are analogous to our modifications in Section 3, which are as follows: first, we assume G, H, and G_T all have two subgroups, and we initially operate solely in the first of these subgroups, so that $\mathcal{A}$ is given $(g_1^{R(x_1,\ldots,x_c)}, h_1^{S(x_1,\ldots,x_c)}, e(g_1,h_1)^{T(x_1,\ldots,x_c)})$ rather than values in the full group. Second, we again switch from the canonical generators g_1 and h_1 to random generators u_1 and v_1. To make our proofs cleaner, we phrase this requirement as follows: for every $\rho_i \in R$, there must exist an efficiently computable function $\hat{\rho}_i$ such that $g_1^{\rho_i(\vec{x})} = u_1^{\hat{\rho}_i(\vec{x})}$, and there must also exist an efficiently computable function $\hat{f}$ such that $g_1^{f(\vec{x})} = u_1^{\hat{f}(\vec{x})}$. Practically, suppose that $u_1 = g_1^r$. Then, using $x_1 = r$, our requirement is equivalent to the requirement that $\rho_i(x_1,\ldots,x_c) = x_1 \cdot \hat{\rho}_i(x_2,\ldots,x_c)$ (and the same for f). Again, we stress that this is just a base translation rather than a restriction on the parameters of the uber-assumption.

Theorem 4.2. *For a bilinear group $\mathbb{G} = (N, G, H, G_T, e, g_1, g_2) \in [\mathsf{BilinearGen}(1^\lambda, 2)]$, consider the decisional uber-assumption parameterized by c, $R = \langle 1, \rho_1(\vec{x}_1), \ldots, \rho_r(\vec{x}_1)\rangle$, $S = T = \langle 1\rangle$, and $f(\vec{x}_1)$. Then, if subgroup hiding holds in $\mathbb{G}$ with respect to $\mu = \{g_1, g_2\}$ and parameter hiding holds with respect to $R \cup \{f\}$, this assumption is implied by the decisional uber-assumption parameterized by ℓc, $R' = \langle 1, \sum_{i=1}^{\ell} \rho_1(\vec{x}_i), \ldots, \sum_{i=1}^{\ell} \rho_r(\vec{x}_i)\rangle$, S, T, and $f' = \sum_{i=1}^{\ell} f(\vec{x}_i)$ for all $\ell = poly(\lambda)$.*

A proof of this theorem can be found in the full version of the paper, and also applies when $R = S = \langle 1\rangle$ and only T contains meaningful functions, or more generally in the case when there might not be an efficiently computable pairing. Intuitively, the transitions rely on the same modified dual-system technique that we used in the proof of Theorem 3.1. First, all elements exist only in the G_1 subgroup, operating over the original set of variables $\vec{x}_1$. A shadow copy of these elements is then added into the G_2 subgroup, which goes unnoticed by subgroup hiding. This shadow copy is then switched to operate over a new set of variables $\vec{x}_2$, which is identical by parameter hiding. These new values are then folded

back into the G_1 subgroup, which is again indistinguishable by subgroup hiding. Finally, the G_2 component is eliminated, which is once again indistinguishable by subgroup hiding. The result is now a G_1 component that operates over both $\vec{x}_1$ and $\vec{x}_2$, and the effect is analogous to the extra degree of randomness we obtain in the proof of Theorem 3.1. Repeating this process $\ell - 1$ more times proves the theorem.

To now show why this theorem is useful, we illustrate that the resulting game is often statistically hard, and thus the original uber-assumption is implied solely by subgroup hiding. To start, consider

$$V = \begin{bmatrix} 1 & \rho_1(\vec{x}_1) & \rho_2(\vec{x}_1) & \cdots & \rho_q(\vec{x}_1) & f(\vec{x}_1) \\ 1 & \rho_1(\vec{x}_2) & \rho_2(\vec{x}_2) & \cdots & \rho_q(\vec{x}_2) & f(\vec{x}_2) \\ & & & \ddots & & \\ \vdots & \vdots & & & \vdots & \vdots \\ & & & \ddots & & \\ 1 & \rho_1(\vec{x}_\ell) & \rho_2(\vec{x}_\ell) & \cdots & \rho_q(\vec{x}_\ell) & f(\vec{x}_\ell) \end{bmatrix} \tag{1}$$

We then have the following lemma, which relates the linear independence of the polynomials with the invertibility of the matrix:

Lemma 4.1. *For all $\lambda \in \mathbb{N}$, if the functions in $R \cup \{f\}$ are linearly independent and of maximum degree poly(λ), $\ell = q+2$ for $q =$ poly(λ), and $N = p_1 \cdot \ldots \cdot p_n$ for distinct primes $p_1, \ldots, p_n \in O(2^{poly(\lambda)})$, then with all but negligible probability the matrix V is invertible.*

Proof. If the matrix V is invertible in $\mathbb{Z}/p_i\mathbb{Z}$ for each prime $p_i \mid N$, then it is also invertible in $\mathbb{Z}/N\mathbb{Z}$. To see that V is invertible (with all but negligible probability) in $\mathbb{Z}/p_i\mathbb{Z}$ for all i, define $F = \mathbb{Z}/p_i\mathbb{Z}$ (or, in the case that N is itself prime, define $F = \mathbb{Z}/N\mathbb{Z}$); then V is a matrix over F, where $|F|$ is exponential in λ. If we consider V instead as a matrix over the polynomial ring $F[x_{1,1}, \ldots, x_{1,c}, \ldots, x_{q+2,c}]$, then we can define its determinant to be the polynomial $D(\vec{x}_1, \ldots, \vec{x}_{q+2})$. By the definition of polynomial linear independence, the columns of V are linearly independent, so D is not the zero polynomial.

To consider the linear independence of the matrix over F, we must consider an assignment of concrete values $\vec{a}_1, \ldots, \vec{a}_{q+2}$ for the variables $\vec{x}_1, \ldots, \vec{x}_{q+2}$. To see that $D(\vec{a}_1, \ldots, \vec{a}_{q+2}) \neq 0$ with all but negligible probability — and thus the matrix V is invertible — consider $d = \max_{i=0}^{q}(d_i)$, where $d_0 = \deg(f)$ and $d_i = \deg(\rho_i)$ for all $\rho_i \in R$; then $\deg(D) \leq (q+1)d$. By the Schwartz-Zippel lemma, $\Pr[D(\vec{a}_1, \ldots, \vec{a}_{q+2}) = 0] \leq (q+1)d/|F|$ for $\vec{a}_1, \ldots, \vec{a}_{q+2} \xleftarrow{\$} F$. As $|F|$ is exponential in λ and both q and d are polynomial in λ, the probability is bounded by a negligible function in λ. □

We then have the following corollary, which indicates when we can show that the original decisional assumption is implied by subgroup hiding.

Corollary 4.1. *The decisional uber-assumption parameterized by (c, R, S, T, f) holds with all but negligible probability if (1) subgroup hiding holds in $\mathbb{G}$ with*

respect to $\mu = \{g_1, g_2\}$, (2) parameter hiding holds with respect to $R \cup \{f\}$, (3) $S = T = \langle 1 \rangle$, and (4) the polynomials in $R \cup \{f\}$ are linearly independent.

Proof. By requirements (1), (2), and (3), Theorem 4.2 tells us that the (c, R, S, T, f)-uber assumption is equivalent to the $(\ell c, R', S, T, f')$-uber-assumption. In this latter assumption, the adversary sees values with exponents of the form $\vec{y} = \vec{r} \cdot V$, where $\vec{r}$ is a random vector of length ℓ and V is the $\ell \times (q+2)$ matrix defined in Equation 1. If we use $\ell = q + 2$, then by requirement (4), Lemma 4.1 tells us that V is invertible with all but negligible probability.

We can now use a bijection argument similar to the one in the proof of Theorem 3.1: $\vec{r}$ and $\vec{y}$ are both members of the set S containing all sets of size $q+2$ over $\mathbb{Z}/N\mathbb{Z}$, so multiplication by V maps S to itself. As V is invertible, the map is invertible as well, and is thus a permutation over S. Sampling $\vec{r}$ uniformly at random and then multiplying by V thus yields a vector $\vec{y}$ that is distributed uniformly at random over $\mathbb{Z}/N\mathbb{Z}$.

An adversary $\mathcal{A}$ thus has no advantage in distinguishing between $\vec{y}$ and a uniformly random vector in S, as the distributions over the two are identical, and thus has no advantage in d-UBER$^{\mathcal{A}}_{\ell c, R', S, T, f'}(\lambda)$. □

As observed by Boneh, Boyen, and Goh, if f is not linearly independent from all polynomials in $R \cup T$, then the assumption becomes trivially false. It furthermore unnecessarily expands the size of the tuple to use polynomials in R or T that are linearly dependent, as, e.g., g^{2x} is redundant given g^x. We therefore believe that the requirement that the polynomials in $R \cup T \cup \{f\}$ be linearly independent is not restrictive, and in fact — to the best of our knowledge — it is satisfied by all existing instantiations of the uber-assumption.

As a concrete example, we finally examine the exponent q-SDH assumption, as introduced and used by Zhang et al. [**?**].

Example 4.1. For exponent q-SDH, $R = \langle 1, \alpha, \alpha^2, \ldots, \alpha^q \rangle$ and $f(\alpha) = \alpha^{q+1}$. Plugging these values into the matrix V gives a Vandermonde matrix, which is invertible. By Corollary 4.1, exponent q-SDH is thus implied by subgroup hiding, assuming parameter hiding holds with respect to the set $\{f_k(\alpha) = \alpha^k\}_{k=1}^{q+1}$ (which, given our discussion in Section 5, currently restricts us to composite-order groups).

4.3 A second approach: computational assumptions in the source group

Although our results in the previous section have potentially broad implications, the requirements for Theorem 4.2 — and in particular the requirement that $S = \langle 1 \rangle$ — are somewhat restrictive, as many q-type assumptions require meaningful functions on both sides of the pairing. We furthermore do not seem able to relax this requirement using our current proof strategy: briefly, the fact that we need subgroup hiding between both G_1 and $G_1 \times G_2$ and between $G_1 \times G_2$ and G_2 means that we cannot give out the subgroup generators h_1 and h_2 on the other

side of the pairing. To get around this restriction and allow meaningful functions on both sides of the pairing, we now consider an alternate approach in which we require subgroup hiding only between G_1 and $G_1 \times G_2$, which allows us to give out h_1.

Theorem 4.3. *For a bilinear group $\mathbb{G} = (N, G, H, G_T, e) \in [\mathsf{BilinearGen}(1^\lambda, 2)]$, consider the computational uber-assumption parameterized by c, $R = \langle 1, \rho_1(\vec{x}), \ldots, \rho_r(\vec{x})\rangle$, S, T, and f. Then, if subgroup hiding holds in $\mathbb{G}$ with respect to $\mu = \{g_1, g_2, h_1\}$ and parameter hiding holds with respect to $R \cup \{f\}$, this is implied by the following assumption for all $\ell = poly(\lambda)$: given*

$$(\mathbb{G}, u_1 g_2^{\sum_{i=1}^{\ell} r_i}, \{u_1^{\rho_k(\vec{x})} g_2^{\sum_{i=1}^{\ell} r_i \rho_k(\vec{x}_i)}\}_{k=1}^{r}, v_1^{S(x_1,\ldots,x_c)}, e(u_1, v_1)^{T(x_1,\ldots,x_c)})$$

for $\vec{x}, r_1, \vec{x}_1, \ldots, r_\ell, \vec{x}_\ell \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$, it is difficult to compute $u_1^{f(\vec{x})} g_2^{\sum_{i=1}^{\ell} r_i f(\vec{x}_i)}$.

A proof of this theorem can be found in the full version of the paper. Intuitively, the starting point is the same as in our previous proofs: all elements exist only in the G_1 subgroup, operating over the original set of variables $\vec{x}$, and a shadow copy of these elements is added into the G_2 subgroup, which goes unnoticed by subgroup hiding. This shadow copy is then switched to operate over a new set of variables $\vec{x}_1$, which is identical by parameter hiding. Now, rather than attempt to move these new variables back into G_1, we simply repeat the process of adding and re-randomizing the original set of variables into the G_2 subgroup, until we end up with ℓ sets of variables there.

Once again, the usefulness of this theorem is revealed only when we examine what this more complex assumption provides. Interestingly, it is not clear how to show that the decisional assumption holds by a statistical argument, as the isolation of the $\vec{x}$ variables in the G_1 subgroup provides a potentially detectable distribution. Instead, we restrict our attention to computational assumptions in the source group, in which the adversary is required to compute $u_1^{f(\vec{x})} g_2^{\sum_{i=1}^{\ell} r_i f(\vec{x}_i)}$ rather than distinguish it from random. In this setting, we have the following corollary; as its proof is analogous to the proof of Corollary 4.1, we omit it here (but it can be found in the full version of the paper).

Corollary 4.2. *The computational uber-assumption parameterized by (c, R, S, T, f) holds in the source group with all but negligible probability if (1) subgroup hiding holds in $\mathbb{G}$ with respect to $\mu = \{g_1, g_2, h_1\}$, (2) parameter hiding holds with respect to $R \cup \{f\}$, (3) the polynomials in $R \cup \{f\}$ are linearly independent.*

To bring everything together, we examine the q-SDH assumption, as defined by Boneh and Boyen [**?**].

Example 4.2. The q-SDH assumption uses $R = \langle 1, \alpha, \ldots, \alpha^q\rangle$, $S = \langle 1, \alpha\rangle$, $T = \langle 1\rangle$, and asks $\mathcal{A}$ to compute $(c, u^{\frac{1}{\alpha+c}})$. Using Theorem 4.3,[4] this is equivalent (under subgroup and parameter hiding) to an assumption in which $\mathcal{A}$ is

[4] Technically, this assumption doesn't meet the requirements of the theorem, as $\mathcal{A}$ produces a new value c rather than a function $f(\vec{x})$. The proof of the theorem can, however, be trivially extended to support assumptions of this type as well, as long as the group satisfies adaptive parameter hiding.

given $(u_1 g_2^{\sum_{i=1}^{q+2} r_i}, u_1^{\alpha} g_2^{\sum_{i=1}^{q+2} r_i \gamma_i}, \ldots, u_1^{\alpha^q} g_2^{\sum_{i=1}^{q+2} r_i \gamma_i^q}, v_1, v_1^{\alpha})$, where $\gamma_1, \ldots, \gamma_{q+2} \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$, and is asked to compute $(c, u_1^{\frac{1}{\alpha+c}} g_2^{\sum_i \frac{r_i}{\gamma_i+c}})$. Applying the same analysis as above, we can ignore G_1 and focus on G_2, in which we use the matrix

$$A = \begin{bmatrix} 1 & \gamma_1 & \cdots & \gamma_1^q & \frac{1}{\gamma_1+c} \\ 1 & \gamma_2 & \cdots & \gamma_2^q & \frac{1}{\gamma_2+c} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & \gamma_\ell & \cdots & \gamma_\ell^q & \frac{1}{\gamma_\ell+c} \end{bmatrix}$$

Then, for its choice of c, $\mathcal{A}$ is given the first $q+1$ entries of $\vec{r} \cdot A$ and needs to compute the final entry. This matrix is invertible, so the same bijection argument as in Corollary 4.2 thus implies that $\mathcal{A}$ can produce the correct value with at most negligible probability, which implies (assuming parameter hiding holds with respect to $\{\rho_k(\alpha) = \alpha^k\}_{k=1}^q$) that q-SDH is implied by subgroup hiding.

5 Instantiating Our Results

Abstractly, our results provide quite a strong guarantee: as long as subgroup hiding and parameter hiding hold, many instantiations of the uber-assumption hold (as well as non-uber-assumptions, such as q-SDH), as they reduce to assumptions that hold by a statistical argument. Concretely, we need to examine which groups support these underlying assumptions.

Parameter hiding. Our strongest requirement in our analysis was the generality of parameter hiding: to reason about any q-type assumption, we need a group where parameter hiding holds for all rational functions. While this seems hard to achieve in general, it does hold for any composite-order group (e.g., any group of order $N = p_1 p_2$ for primes p_1 and p_2), as the value of any exponent modulo p_1 is independent of its value modulo p_2.

Subgroup hiding. In groups without a pairing — such as the target group of a bilinear tuple or a group over a non-pairing-friendly elliptic curve — subgroup decision is fairly straightforward. In groups with a pairing, however, the concerns mentioned in Section 2 (in which certain subgroup generators on the other side of the pairing could render subgroup decision easy) mean we have to be more careful. Our first approach in Section 4.2 relies on being unable to distinguish random elements of both G_1 and G_2 from $G_1 \times G_2$, even when given g_1 and g_2. This cannot hold, for example, in a symmetric bilinear group, so this assumption is reasonable only in the asymmetric setting. Our second approach in Section 4.3 requires that subgroup hiding holds even given h_1 and g_2, so it again requires an asymmetric pairing.

Instantiations. As mentioned above, our results in Sections 3 and 4 can be applied in any composite-order group where we can assume subgroup hiding.

Reasonable candidates for such a group include composite-order elliptic curve groups without efficient pairings, the target group of a composite-order bilinear group, or composite-order subgroups of finite fields.

In the case where we do have a pairing, we need an asymmetric composite-order bilinear group in order to make subgroup hiding a reasonable assumption. Although most composite-order bilinear groups are symmetric (as they are groups of points on supersingular curves), ordinary composite-order curves were first introduced by Boneh, Rubin, and Silverberg [**?**], and their applicability for cryptography — and in particular an examination of the nature of the resulting asymmetric composite-order bilinear group — was very recently explored by Meiklejohn and Shacham [**?**].

Applications. In asymmetric composite-order bilinear groups we can prove a wide range of constructions secure based on just subgroup hiding. For example, our examination of q-SDH means that the Boneh-Boyen signature, the Boneh-Boyen-Shacham group signature [**?**], and the attribute-based signature due to Maji et al. [**?**] can all be proved secure under subgroup hiding, and the fact that q-DHI [**?**] is also equivalent to subgroup hiding implies the Dodis-Yampolskiy VUF and the Jarecki-Liu PRF [**?**] can also both be proved secure based on subgroup hiding.

6 Conclusions and Open Problems

This paper demonstrated the applicability of the dual-system technique (and variants on it) by first proving the security of the Dodis-Yampolskiy PRF — using a domain of arbitrary size — under subgroup hiding, and then proving equivalence between many classes of the uber-assumption. This latter result further implies that many of these classes are in fact implied solely by subgroup hiding, as they reduce to assumptions that hold by a purely statistical argument. Our paper thus demonstrates that many common q-type assumptions — and the constructions that rely on them for security — can be implied directly by subgroup hiding when instantiated in the appropriate bilinear groups.

As our paper is a first step, many interesting directions and open problems remain. For example, we currently cannot prove anything about, e.g., decisional assumptions — such as q-DDHE — that require meaningful functions on both sides of the pairing. Perhaps the biggest open problem is obtaining more robust forms of parameter hiding in prime-order groups. Prime-order groups have the benefit of being significantly more efficient, and it is possible to construct groups with the appropriate subgroup hiding requirements using dual pairing vector spaces [**?**,**?**], as exemplified most recently by Lewko and Meiklejohn [**?**].

For parameter hiding in prime-order bilinear groups, however, it is currently known how to obtain parameter hiding only for linear functions. Papers that have focused on translating these structural properties into prime-order settings, however, have indicated that they focus on such simple functions to keep their "constructions... simple and tailored to the requirements that [they] need" [**?**],

so we consider constructing parameter hiding for more robust functions in the prime-order setting an interesting open problem rather than an impossibility.

Acknowledgments

We thank Michael Naehrig for his valuable feedback, and the anonymous reviewers for their helpful suggestions.