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

Melissa Chase¹ and Sarah Meiklejohn^{2,*}

¹ Microsoft Research Redmond, USA

`melissac@microsoft.com`

² UC San Diego, USA

`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 [10,3,5,35], group signatures [7,13,20], zero-knowledge proofs [21,22], (hierarchical) identity-based encryption [8,4,6,31], and functional and attribute-based encryption [29,33,34]. 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) [8] and Decision Linear [7] 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 $(g, g^x, g^{x^2}, \dots, g^{x^q})$ and is asked to produce $e(g, g)^{1/x}$. While the ‘‘uber-assumption’’ [6,12] generalizes many q -type assumptions (as well as many static assumptions) and provides a lower bound for their security in the generic group model [43], 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 [17] 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 [24] showed that Boneh-Boyen signatures were in fact equivalent to the q -SDH assumption that they rely on for security. Finally, Cheon [16] 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 $(g, g^x, \dots, g^{x^q})$). 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 [44] in 2009 and has been used extensively since [31,29,30,33,28,34]. Briefly, this technique takes advantage of *subgroup hiding* in bilinear groups [9]; 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* [28], 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 [14,2,15,26].

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 [11,37].

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),

subgroup 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 [44], and was applied subsequently to achieve a wide variety of results [31,29,41,30,33,32,28,42], 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 [6] 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 [23], 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. Anecdotaly, several results use the dual-system technique to eliminate the requirement on q -type assumptions for specific primitives or constructions: Gerbush et al. [18] 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 [1]; and the original result of Waters [44] 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 $y \leftarrow A(x_1, \dots, x_n; R)$ we denote running algorithm A on inputs $x_1, \dots, x_n$ and random coins R and assigning its output to y . By $y \xleftarrow{\$} A(x_1, \dots, x_n)$ we denote $y \leftarrow A(x_1, \dots, x_n; R)$ for some random coins R . By $[A(x_1, \dots, x_n)]$ we denote the set of values that have positive probability of being output by A on inputs $x_1, \dots, x_n$. 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. $\text{Pr}[G]$ 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 $e : G \times H \rightarrow G_T$ 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 **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 [9]; more general definitions exist as well [28,27].)

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 $\mathbf{Adv}_{\mathcal{A}}^{sgh}(\lambda) = 2\Pr[SGH_{\mu}^{\mathcal{A}}(\lambda)] - 1$, where $SGH_{\mu}^{\mathcal{A}}(\lambda)$ is defined as follows for $\mu \subseteq \{g_1, g_2, h_1, h_2\}$:*

$\text{MAIN } SGH_{\mu}^{\mathcal{A}}(\lambda)$
 $b \xleftarrow{\$} \{0, 1\}; (N, G, H, G_T, e, g_1, g_2, h_1, h_2) \xleftarrow{\$} \text{BilinearGen}(1^{\lambda}, 2)$
 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 $\mathbf{Adv}_{\mathcal{A}}^{sgh}(\lambda) < \nu(\lambda)$.

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 [28] 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 [27]; 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 $\{g_1^{f(x_1, \dots, x_n)} g_2^{f(x_1, \dots, x_n)}\}_{f \in \mathcal{F}}$ is identical to $\{g_1^{f(x_1, \dots, x_n)} g_2^{f(x'_1, \dots, x'_n)}\}_{f \in \mathcal{F}}$ for $g_1 \xleftarrow{\$} G_1$, $g_2 \xleftarrow{\$} G_2$, and $x_1, x'_1, \dots, x_n, x'_n \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$. (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 $\text{poly}(\lambda)$, the distribution $\{g_1^{f(x_1, \dots, x_n)} g_2^{f(x_1, \dots, x_n)}\}_{f \in S}$ is statistically close to $\{g_1^{f(x_1, \dots, x_n)} g_2^{f(x'_1, \dots, x'_n)}\}_{f \in S}$ for $g_1 \xleftarrow{\$} G_1$, $g_2 \xleftarrow{\$} G_2$, and $x_1, x'_1, \dots, x_n, x'_n \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$.*

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 [19] $\mathcal{F}$ specifies the algorithms $\mathcal{F}.\text{Pg}$, $\mathcal{F}.\text{Keys}$, $\mathcal{F}.\text{Dom}$, $\mathcal{F}.\text{Rng}$, and $\mathcal{F}.\text{Ev}$. Via $fp \xleftarrow{\$} \mathcal{F}.\text{Pg}(1^\lambda)$ one generates a description fp of a function $\mathcal{F}.\text{Ev}(1^\lambda, fp): \mathcal{F}.\text{Keys}(1^\lambda, fp) \times \mathcal{F}.\text{Dom}(1^\lambda, fp) \rightarrow \mathcal{F}.\text{Rng}(1^\lambda, fp)$. The evaluation algorithm $\mathcal{F}.\text{Ev}$ is PT and deterministic.

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

<u>MAIN $\text{PRF}_F^{\mathcal{A}}(\lambda)$</u>	<u>Procedure $\text{FN}_{sk}(x)$</u>
$b \xleftarrow{\$} \{0, 1\}; fp \xleftarrow{\$} F.\text{Pg}(1^\lambda); sk \xleftarrow{\$} F.\text{Keys}(1^\lambda, fp)$	if $b = 0$ $y \xleftarrow{\$} F.\text{Rng}(1^\lambda, fp)$
$b' \xleftarrow{\$} \mathcal{A}^{\text{FN}}(1^\lambda, fp)$	if $b = 1$ $y \leftarrow F.\text{Ev}(1^\lambda, fp, sk, x)$
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 $\mathbf{Adv}_{F,\mathcal{A}}^{\text{prf}}(\lambda) \leq \nu(\lambda)$.

3 Pseudorandom Functions

In this section, we explore the security of the Dodis-Yampolskiy PRF [17]. First, we recall the Dodis-Yampolskiy PRF, instantiated for our purposes in a group of composite order $N = p_1 p_2$.¹

- $F.\text{Pg}(1^\lambda)$: Output $(N, G, H, G_T, e, g, h) \xleftarrow{\$} \text{BilinearGen}(1^\lambda, 2)$. Then $F.\text{Keys} = F.\text{Dom} = \mathbb{Z}/N\mathbb{Z}$, and $F.\text{Rng} = G_T$.
- $F.\text{Ev}(1^\lambda, fp, sk, x)$: 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 $(g, g^x, \dots, g^{x^q})$, 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 $\mathbf{Adv}_{F,\mathcal{A}}^{\text{prf}}(\lambda) \leq 2^{a(\lambda)} \cdot \mathbf{Adv}_{\mathcal{A}}^{2^{a(\lambda)}\text{-DBDHI}}(\lambda)$, 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

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

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

¹ 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 $fp \in [\mathbb{F}.\text{Pg}(1^\lambda)]$, if subgroup hiding holds in $\mathbb{G}$ and adaptive parameter hiding holds with respect to $\{\mathbb{F}.\text{Ev}(1^\lambda, fp, \cdot, x)\}_{x \in \mathbb{F}.\text{Dom}}$, then $\mathbb{F}$ 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” [6,12], 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 $R = \langle \rho_1(x_1, \dots, x_c), \dots, \rho_r(x_1, \dots, x_c) \rangle$ and as shorthand use $\rho_i(\vec{x}) = \rho_i(x_1, \dots, x_c)$ and $g^{R(x_1, \dots, x_c)} = \{g^{\rho_i(\vec{x})}\}_{i=1}^r$ (and similarly for S and T).

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

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

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

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 $\text{d-UBER}_{c,R,S,T,f}^{\mathcal{A}}(\lambda)$ for the decisional uber-assumption is defined analogously, except rather than compute $g^{f(x_1, \dots, x_c)}$ 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 $(g_1^{R(x_1, \dots, x_c)}, h_1^{S(x_1, \dots, x_c)}, e(g_1, h_1)^{T(x_1, \dots, x_c)})$ 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 $\rho_i(x_1, \dots, x_c) = x_1 \cdot \hat{\rho}_i(x_2, \dots, x_c)$ (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 $\mathbb{G} = (N, G, H, G_T, e, g_1, g_2) \in [\text{BilinearGen}(1^\lambda, 2)]$, consider the decisional uber-assumption parameterized by c , $R = \langle 1, \rho_1(\vec{x}_1), \dots, \rho_r(\vec{x}_1) \rangle$, $S = T = \langle 1 \rangle$, 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 , $R' = \langle 1, \sum_{i=1}^\ell \rho_1(\vec{x}_i), \dots, \sum_{i=1}^\ell \rho_r(\vec{x}_i) \rangle$, S , T , and $f' = \sum_{i=1}^\ell f(\vec{x}_i)$ for all $\ell = \text{poly}(\lambda)$.*

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

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

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 $\text{poly}(\lambda)$, $\ell = q+2$ for $q = \text{poly}(\lambda)$, and $N = p_1 \cdots p_n$ for distinct primes $p_1, \dots, p_n \in O(2^{\text{poly}(\lambda)})$, 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 $F[x_{1,1}, \dots, x_{1,c}, \dots, x_{q+2,c}]$, then we can define its determinant to be the polynomial $D(\vec{x}_1, \dots, \vec{x}_{q+2})$. 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 $\vec{a}_1, \dots, \vec{a}_{q+2}$ for the variables $\vec{x}_1, \dots, \vec{x}_{q+2}$. To see that $D(\vec{a}_1, \dots, \vec{a}_{q+2}) \neq 0$ 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, $\Pr[D(\vec{a}_1, \dots, \vec{a}_{q+2}) = 0] \leq (q+1)d/|F|$ for $\vec{a}_1, \dots, \vec{a}_{q+2} \xleftarrow{\$} F$. 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 . $\square$

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 $\text{d-UBER}_{\ell c, R', S, T, f'}^A(\lambda)$. $\square$

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. [45].

Example 4.1. For exponent q -SDH, $R = \langle 1, \alpha, \alpha^2, \dots, \alpha^q \rangle$ 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 $\mathbb{G} = (N, G, H, G_T, e) \in [\text{BilinearGen}(1^\lambda, 2)]$, consider the computational uber-assumption parameterized by $c, R = \langle 1, \rho_1(\vec{x}), \dots, \rho_r(\vec{x}) \rangle, S, T$, 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 $\ell = \text{poly}(\lambda)$: given*

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

for $\vec{x}, r_1, \vec{x}_1, \dots, r_\ell, \vec{x}_\ell \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$, 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 [5].

Example 4.2. The q -SDH assumption uses $R = \langle 1, \alpha, \dots, \alpha^q \rangle$, $S = \langle 1, \alpha \rangle$, $T = \langle 1 \rangle$, and asks $\mathcal{A}$ to compute $(c, u^{\frac{1}{\alpha+c}})$. Using Theorem 4.3,² this is equivalent (under subgroup and parameter hiding) to an assumption in which $\mathcal{A}$ is

² 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 $(u_1 g_2^{\sum_{i=1}^{q+2} r_i}, u_1^\alpha g_2^{\sum_{i=1}^{q+2} r_i \gamma_i}, \dots, u_1^{\alpha^q} g_2^{\sum_{i=1}^{q+2} r_i \gamma_i^q}, v_1, v_1^\alpha)$, where $\gamma_1, \dots, \gamma_{q+2} \xleftarrow{\$} \mathbb{Z}/N\mathbb{Z}$, 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 [11], 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 [37].

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 [7], and the attribute-based signature due to Maji et al. [36] can all be proved secure under subgroup hiding, and the fact that q -DHI [38] is also equivalent to subgroup hiding implies the Dodis-Yampolskiy VUF and the Jarecki-Liu PRF [25] 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 [39,40], as exemplified most recently by Lewko and Meiklejohn [27].

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” [27],

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.

References

1. Attrapadung, N., Libert, B.: Functional encryption for inner product: Achieving constant-size ciphertexts with adaptive security or support for negation. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 384–402. Springer, Heidelberg (2010)
2. Au, M.H., Susilo, W., Mu, Y.: Constant-size dynamic k -TAA. In: De Prisco, R., Yung, M. (eds.) SCN 2006. LNCS, vol. 4116, pp. 111–125. Springer, Heidelberg (2006)
3. Boldyreva, A.: Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In: Desmedt, Y.G. (ed.) PKC 2003. LNCS, vol. 2567, pp. 31–46. Springer, Heidelberg (2003)
4. Boneh, D., Boyen, X.: Secure identity based encryption without random oracles. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 443–459. Springer, Heidelberg (2004)
5. Boneh, D., Boyen, X.: Short signatures without random oracles. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 56–73. Springer, Heidelberg (2004)
6. Boneh, D., Boyen, X., Goh, E.-J.: Hierarchical identity based encryption with constant size ciphertext. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 440–456. Springer, Heidelberg (2005)
7. Boneh, D., Boyen, X., Shacham, H.: Short group signatures. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 41–55. Springer, Heidelberg (2004)
8. Boneh, D., Franklin, M.: Identity-based encryption from the Weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
9. Boneh, D., Goh, E.-J., Nissim, K.: Evaluating 2-DNF formulas on ciphertexts. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 325–341. Springer, Heidelberg (2005)
10. Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 514–532. Springer, Heidelberg (2001)
11. Boneh, D., Rubin, K., Silverberg, A.: Finding ordinary composite order elliptic curves using the Cocks-Pinch method. *Journal of Number Theory* 131(5), 832–841 (2011)
12. Boyen, X.: The uber-assumption family (invited talk). In: Galbraith, S.D., Paterson, K.G. (eds.) Pairing 2008. LNCS, vol. 5209, pp. 39–56. Springer, Heidelberg (2008)
13. Boyen, X., Waters, B.: Compact group signatures without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 427–444. Springer, Heidelberg (2006)
14. Camenisch, J., Hohenberger, S., Lysyanskaya, A.: Compact e-cash. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 302–321. Springer, Heidelberg (2005)

15. Camenisch, J., Lysyanskaya, A., Meyerovich, M.: Endorsed e-cash. In: 2007 IEEE Symposium on Security and Privacy, pp. 101–115. IEEE Computer Society Press (May 2007)
16. Cheon, J.H.: Security analysis of the strong Diffie-Hellman problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 1–11. Springer, Heidelberg (2006)
17. Dodis, Y., Yampolskiy, A.: A verifiable random function with short proofs and keys. In: Vaudenay, S. (ed.) PKC 2005. LNCS, vol. 3386, pp. 416–431. Springer, Heidelberg (2005)
18. Gerbush, M., Lewko, A., O'Neill, A., Waters, B.: Dual form signatures: An approach for proving security from static assumptions. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 25–42. Springer, Heidelberg (2012)
19. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. In: 25th FOCS, pp. 464–479. IEEE Computer Society Press (October 1984)
20. Groth, J.: Simulation-sound NIZK proofs for a practical language and constant size group signatures. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 444–459. Springer, Heidelberg (2006)
21. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer, Heidelberg (2006)
22. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 415–432. Springer, Heidelberg (2008)
23. Jager, T., Rupp, A.: The semi-generic group model and applications to pairing-based cryptography. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 539–556. Springer, Heidelberg (2010)
24. Jao, D., Yoshida, K.: Boneh-Boyen signatures and the strong Diffie-Hellman problem. In: Shacham, H., Waters, B. (eds.) Pairing 2009. LNCS, vol. 5671, pp. 1–16. Springer, Heidelberg (2009)
25. Jarecki, S., Liu, X.: Efficient oblivious pseudorandom function with applications to adaptive OT and secure computation of set intersection. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 577–594. Springer, Heidelberg (2009)
26. Lee, M.Z., Dunn, A.M., Katz, J., Waters, B., Witchel, E.: Anon-pass: practical anonymous subscriptions. In: Proceedings of IEEE Symposium on Security and Privacy (2013)
27. Lewko, A., Meiklejohn, S.: A profitable sub-prime loan: Obtaining the advantages of composite-order in prime-order bilinear groups. Cryptology ePrint Archive, Report 2013/300 (2013), <http://eprint.iacr.org/2013/300>
28. Lewko, A.: Tools for simulating features of composite order bilinear groups in the prime order setting. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 318–335. Springer, Heidelberg (2012)
29. Lewko, A., Okamoto, T., Sahai, A., Takashima, K., Waters, B.: Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 62–91. Springer, Heidelberg (2010)
30. Lewko, A., Rouselakis, Y., Waters, B.: Achieving leakage resilience through dual system encryption. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 70–88. Springer, Heidelberg (2011)
31. Lewko, A., Waters, B.: New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 455–479. Springer, Heidelberg (2010)

32. Lewko, A., Waters, B.: Decentralizing attribute-based encryption. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 568–588. Springer, Heidelberg (2011)
33. Lewko, A., Waters, B.: Unbounded HIBE and attribute-based encryption. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 547–567. Springer, Heidelberg (2011)
34. Lewko, A., Waters, B.: New proof methods for attribute-based encryption: Achieving full security through selective techniques. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 180–198. Springer, Heidelberg (2012)
35. Lu, S., Ostrovsky, R., Sahai, A., Shacham, H., Waters, B.: Sequential aggregate signatures and multisignatures without random oracles. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 465–485. Springer, Heidelberg (2006)
36. Maji, H.K., Prabhakaran, M., Rosulek, M.: Attribute-based signatures. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 376–392. Springer, Heidelberg (2011)
37. Meiklejohn, S., Shacham, H.: New trapdoor projection maps for composite-order bilinear groups. Cryptology ePrint Archive, Report 2013/657 (2013), <http://eprint.iacr.org/2013/657>
38. Mitsunari, S., Saka, R., Kasahara, M.: A new traitor tracing. IEICE Transactions E85-A(2), 481–484 (2002)
39. Okamoto, T., Takashima, K.: Homomorphic encryption and signatures from vector decomposition. In: Galbraith, S.D., Paterson, K.G. (eds.) Pairing 2008. LNCS, vol. 5209, pp. 57–74. Springer, Heidelberg (2008)
40. Okamoto, T., Takashima, K.: Hierarchical predicate encryption for inner-products. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 214–231. Springer, Heidelberg (2009)
41. Okamoto, T., Takashima, K.: Fully secure functional encryption with general relations from the decisional linear assumption. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 191–208. Springer, Heidelberg (2010)
42. Ramanna, S.C., Chatterjee, S., Sarkar, P.: Variants of Waters’ dual system primitives using asymmetric pairings. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 298–315. Springer, Heidelberg (2012)
43. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 256–266. Springer, Heidelberg (1997)
44. Waters, B.: Dual system encryption: Realizing fully secure IBE and HIBE under simple assumptions. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 619–636. Springer, Heidelberg (2009)
45. Zhang, F., Safavi-Naini, R., Susilo, W.: An efficient signature scheme from bilinear pairings and its applications. In: Bao, F., Deng, R., Zhou, J. (eds.) PKC 2004. LNCS, vol. 2947, pp. 277–290. Springer, Heidelberg (2004)