On Post-Modern Cryptography*

Oded Goldreich†
Department of Computer Science
Weizmann Institute of Science
Rehovot, Israel.
oded.goldreich@weizmann.ac.il

November 20, 2006

It is possible to build a cabin with no foundations, but not a lasting building.

Eng. Isidor Goldreich (1906–1995)

Summary: This essay relates to a recent article of Koblitz & Menezes that "criticizes several typical 'provable security' results" and argues that the "theorem-proof paradigm of theoretical mathematics is often of limited relevance" to cryptography. Although it feels ridiculous to answer such a claim, we undertake to do so in this essay. In particular, we point out some of the fundamental philosophical flaws that underly the said article and some of its misconceptions regarding theoretical research in Cryptography in the last quarter of a century.

Contents

*In order to reflect the philosophical nature of the current essay, we chose to keep away from the common style of scientific articles. Indeed, this essay refers to science, but its methodology cannot possibly be scientific (see last section).

†Provenance: The author has written dozens of articles and a two-volume book that Koblitz & Menezes would have labeled as dealing with "provable security" (and/or as following the "theorem-proof paradigm").

A kind of introduction that does not reveal much[1]

On a nominal level, the title of this essay is justified by the fact that the article of Koblitz & Menezes criticizes a central theme in Modern Cryptography[2] (i.e., rigorous analysis). Thus, one may view the association of their critique with the post-modernist cultural critique as merely a joke. However, we find this association justified for at least one fundamental reason: In our opinion, at the last account, both post-modernism and the "critique of rigorous analysis in Modern Cryptography" are reactionary (i.e., they play to the hands of the opponents of progress).[3]

We note that, in comparison to the critique of rigorous analysis in Cryptography, there are at least a couple of points in favor of post-modernism. Firstly, post-modernism offers valuable insights on modernity, whereas we fail to identify any such insights in the critique of rigorous analysis in Cryptography. Secondly, by revealing the oppressive character (or potential) of modernity, post-modernism offers a liberating potential. In contrast, the rigorous analysis of cryptography has never gained dominance in any field of power, and thus viewing it as oppressive is quite odd.

A side comment on terminology

As actually stated in the text of Koblitz & Menezes, their critique targets the *rigorous analysis* methodology of cryptography (which evolves around clear definitions and rigorous inference rules). However, most of their text identifies the said methodology with the term "provable security" (which was not invented by Koblitz & Menezes but rather adopted by them). We comment that, within the domain of the rigorous analysis methodology of Cryptography, the term "provable security" is quite odd and rather inappropriate.[4] Let us elaborate.

Security (according to some definition) is a property that some systems may have. In the domain of science, statements are either valid or invalid, and their state of validity can be either known or unknown. Saying that the validity of a statement is known means that the validity can be established based on the accepted methodology of the relevant discipline. Thus, within the domain of a rigorous analysis of cryptography, the term "provable security" (and in general "provable property") makes no sense; that is, the adjective "provable" adds nothing to the claim of security (assuming that the claim is valid, and it cannot be applied if the claim is invalid or unknown to be valid). Indeed, qualifying a noun by an adjective that adds nothing to it is peculiar, but more importantly it may only cause confusion. Specifically, saying that "X is provable secure" suggests that it is legitimate (within the discipline) to claim that "X is secure" without being able to establish this claim by using the methodology that is acceptable in the discipline.[5]

[1]With apologies to Robert Musil's *The Man without Qualities*.

[2]The historical distinction between "Classical Cryptography" and "Modern Cryptography" is irrelevant to most of this essay, because Cryptography is currently associated with Modern Cryptography. Still, in the context of this paragraph this distinction makes sense, because Modern Cryptography refers to the establishment of an academic discipline with a comprehensive research agenda. This event brought about the "rigorous analysis" revolution to which Koblitz & Menezes object.

[3]There is, however, a difference. Extreme post-modernists typically argue for the equality of all perspectives, which in practice plays to the hands of those in power (which, by their reality, are more oppressive than their opponents). Koblitz & Menezes openly argue in favor of the reactionary perspective (which fetishizes intuition).

[4]But, indeed, what Koblitz & Menezes mean to say is that they reject the rigorous analysis methodology of Cryptography. Thus, from their perspective, it makes sense to use the term "provable security" (as shorthand for security as established by this "odd" methodology). However, using this term is inappropriate for somebody who does accept the said methodology and operates within it.

[5]Needless to say, no scientific discipline allows such a situation. In particular, within the domain of the rigorous analysis methodology of cryptography, if one believes that "X is secure" but cannot establish this fact, then one

The adequate methodology for cryptographic research

However, the issue at hand is not a choice of terminology, but rather a choice of methodology; that is, *the issue is a choice of the adequate methodology for cryptographic research*. Specifically, the question is *whether cryptographic research should adhere to the rigorous analysis methodology*. One may dismiss this question by saying that cryptographic research is part of computer science, which in turn is part of science, where the latter is the domain of rigorous analysis. While this answer may satisfy many readers, philosophically-inclined readers may (rightfully) ask *why should cryptographic research be part of science*. Indeed, the text of Koblitz & Menezes does raise this question.

Before addressing this question, let us note that *cryptographic research is indeed part of science*.[6] This assertion is empirical and it refers to the current sociology of the discipline; that is, we believe that a vast majority of the members of this research community identify themselves as scientists. This is the reason that most of these researchers may dismiss the foregoing question. Still, one may ask (as Koblitz & Menezes seem to do) whether the research community is wrong in identifying itself as scientific and/or whether it should change its research methodology.

This question may seem irrelevant to researchers who view themselves firstly as scientists and secondly as scientists that specialize in cryptography.[7] For such researchers, the commitment to the rigorous analysis methodology comes before the commitment to cryptographic research. However, the question remains valid: it does not refer to personal choices of individual members of the cryptographic research community, but rather to a hypothetical choice of the discipline itself. While personal choices of individuals may be based on various considerations (e.g., a primary commitment to Science at large), the choice of the discipline itself should be based on its intrinsic logic as reflected in it founding questions.

Being done with the preliminary clarifications regarding the nature of the question at stake, we now turn to the answer. In our opinion, the answer to the question of whether cryptographic research should be committed to rigorous analysis is a big YES. *In general, we believe that rigorous analysis is, by far, the best way to study reality.*[8] *Moreover, in the case of cryptography, this general principle is more important than in any other discipline.*

The foregoing assertion is based on the realization that cryptography is focused on adversarial behavior; that is, the protection against adversarial behavior is the discipline's founding question. Needless to say, adversarial behavior is very different from normal behavior. Furthermore, it is almost always the case that the (adversarial) behavior that harms a system is of a type that the system's designer did not expect. In contrast, most disciplines are concerned with normal behavior, or with deviations from the norm that one has already observed or can envision. Our point is that, *while a rigorous analysis is of great value for questions regarding normal behavior, it is indispensable for questions regarding abnormal and unexpected behavior*. Let us elaborate.

The design of cryptographic systems is a very difficult task. One cannot rely on intuitions regarding the typical state of the environment in which the system operates. For sure, the *adversary* attacking the system will try to manipulate the environment into untypical states. Nor can one be content with counter-measures designed to withstand specific attacks, since the adversary (which acts after the design of the system is completed) will try to attack the system in ways that are typically different from the ones the designer had envisioned. The validity of the foregoing assertions seems self-evident, still some people hope that in practice ignoring these tautologies will not result

should state "X is secure" as a conjecture.

[6]N.B., we refer to cryptographic research and not to other activities that may be viewed as related to cryptography.

[7]Indeed, the author of this essay views himself in this way.

[8]We refrain from justifying this opinion, which is a central pivot of modernity and has been the subject of numerous philosophical works (starting, say, with Sir Francis Bacon's *Novum Organum*).

in actual damage. Experience shows that these hopes rarely come true; cryptographic systems based on make-believe are broken, typically sooner than later.

In view of the foregoing, we believe that it makes little sense to make assumptions regarding the specific *strategy* that the adversary may use. The only assumptions that can be justified refer to the computational *abilities* of the adversary. Furthermore, it is our opinion that the design of cryptographic systems has to be based on firm foundations, whereas *ad hoc* approaches and heuristics are a very dangerous way to go. A heuristic may make sense when the designer has a very good idea about the environment in which a system is to operate, yet a cryptographic system has to operate in a maliciously selected environment which typically transcends the designer's view. This situation calls for adopting a rigorous analysis of security, which is based on clear definitions and rigorous inference rules.

Let us spell-out the dangers in the alternative of using vague specifications and/or inference rules that are only supported by intuition. While vague specifications are always bad practice, their harmfulness with respect to adversarial behavior cannot be overstated. Failure to specify one's security concerns is most likely to lead to the realization that these concerns were violated by a clever adversary. Indeed, in retrospect one can identify damage cause by an adversary, and this damage can be described by using intuitive and vague language. But the goal of a cryptographic system is never to get to the situation of describing damage caused by an adversary, and specifying what this means (*a priori*) cannot be done by using vague language. Likewise, inference rules that are only supported by intuition are inadequate, because the intuition used in the argument is the one of the designer while the intuition that really counts is the one of the adversary (let alone that the adversary acts after the designer has explicitly or implicitly made its intuitive claims). In contrast, a rigorous inference is universally valid.

Needless to say, Koblitz & Menezes do not advocate a full rejection of rigorous analysis. It is only that they lack a commitment to this methodology: they are willing to apply it when it suits them and feel free to ignore it otherwise. Specifically, their text represents the approach that feelings and intuitions (which lack any real justification) are superior to knowledge obtained via rigorous analysis *in the sense that whenever the two disagree they prefer the former.*[9] In our opinion, this approach is extremely dangerous in the context of cryptographic research and is utterly unscientific in general. Regarding the specific context of cryptography, we have already discussed the danger of relying on intuition in matters that clearly transcend intuition (i.e., adversarial behavior). Regarding the general attitude of Science towards intuition, it is of keen interest but not of trust. (We conclude this section with a short comment on this issue and return to it in the last section.)

Science values intuitive ideas and seeks to explore their validity. Presented with an intuition (regarding some topic), Science hopes to put this intuition on sound grounds (or modify it such that the modified/qualified intuition can be rigorously justified), but is willing (and forced) to abandon it if proved false. At a last resort, Science stays ignorant with respect to intuition, hoping to redeem this our state of affairs in the future. However, Science always stays committed to its own methodology. It is also not afraid to make conjectures, but it keeps a clear distinction between these and facts.

[9] Needless to say, the fact that Koblitz & Menezes are apparently willing to apply a rigorous analysis whenever it yields a result that fits their intuition is of little significance. The question is what do they prefer when the rigorous analysis disagrees with their intuition.

Applied science: theory versus practice

Before continuing, let us clarify that this essay is focused on cryptographic research, and not on the application of this research to practice. Still a few comments regarding the latter issue are in place.

The general principle that governs the application of theoretical research to practice is that (scientific) research *informs* (technological) practice. This does not mean that practice reduces to a straightforward implementation of theoretical results. On the contrary, the application of theoretical results in practice requires a deep (but not necessarily detailed) understanding of theory as well as the exercising of judgment (which in turn is based on the principles that underly the theory). Indeed, while one can make important research contributions without having a deep understanding of the principles that underly the theory, it seems much harder to design a good practical system without such an understanding.

In particular, in our opinion, the principles that underly the theory of cryptography are the focus on clear definitions of security and the application of rigorous inferences regarding security. Thus, we believe that practice should be based on three ingredients: (1) using clear definitions of the one's goals, (2) using clear definitions of one's assumptions, and (3) providing a rigorous justification of the claim that if the stated assumptions hold then the designed system meets the stated goals.

We shall return to the first ingredient later in this essay. Regarding the second ingredient, note that we are not expressing an opinion on which assumptions to use, except that we insist that they be clear. The clarity of the assumptions is positively correlated to their simplicity. The simpler the assumptions, the better estimate we may have regarding their validity. Indeed, this is the reason that one should prefer assumptions regarding the intractability of some simple tasks (e.g., inverting a one-way function) over the assumption that the designed system satisfies the relevant (cryptographic) specifications. The point being that the latter specification and certainly the designed system are typically too complex to allow for an intuitive evaluation, let alone that it may be the case that the specification is self-contradictory (and cannot be met at all). Thus, it is advised to use significantly simpler assumptions, and to provide a rigorous analysis relating these assumptions to the claim that the designed system meets the specification.

More on assumptions: inequality and choice

The last paragraph touched on the devil's argument by which, since we are using assumptions anyhow, why don't we just assume that the designed system meets the postulated specifications. As already stated, our view is that not all assumptions are equal. Specifically, we distinguish assumptions by their clarity and simplicity, and argue that the validity of clear and simpler assumptions is easier to evaluate.[10] Let us demonstrate the point with a few examples.

Consider, for example, the definition of one-way functions and the definition of zero-knowledge interactive proofs (ZKIPs). While the definition of one-way functions refers to the intractability of a standard computational task, the definition of ZKIPs refers to a significantly more complex situation. Specifically, the latter definition refers to two interactive machines and to two seemingly conflicting requirements, representing the security concerns of each of the two parties.[11] On one

[10]We mention that, a few years ago, Moni Naor suggested to classify assumptions according to the complexity of verifying counterexamples to them.

[11]Indeed, these intuitively conflicting requirements yield an impossibility result when uni-directional communication is concern. However, the intuition by which this conflict yield an impossibility result fails when one allows more

hand, it is required that the verifier's strategy protects against any adversarial attempt to fool the verifier into accepting false assertions. On the other hand, it is required that the prover's strategy yields no knowledge or rather protects against any adversarial attempt to extract knowledge out of the prover. The latter requirement is formulated using the simulation paradigm, which is quite non-trivial by itself, and in the case of computational zero-knowledge this means that the simulation is indistinguishable from the real interaction by yet a third adversarial entity (i.e., the distinguisher). Things become even more complex when considering the definition of secure multi-party computation.[12] Furthermore, even in the case of simpler systems, such as encryption schemes, the security definition is quite non-trivial (and fairly complex relative to the definition of one-way functions); for example, consider the definition of (semantic) security for encryption schemes (even in the passive model, and moreover when wishing to protect against various types of active attacks).

Indeed, when the definition of zero-knowledge interactive proofs was first put forward, the question of its viability was far from being clear; that is, researchers did wonder whether there exists any hard problems that have zero-knowledge interactive proofs. Furthermore, the same reaction arises in class whenever one teaches the subject. In contrast, the reactions (were and) are very different when the definition of one-way functions is concerned, and this is due to the simplicity of the latter definition and to its relation to very familiar phenomena (i.e., the existence of processes that seem hard to reverse). Thus, it is of great interest to note that if one-way function exist then there exists hard problems that have zero-knowledge interactive proofs.

We stress that even today, twenty years after the foregoing result was established, we know of no intuitive reason to suspect that there exists hard problems that have zero-knowledge interactive proofs. On the contrary, the uninformed intuition would suggest the opposite (i.e., that no such interactive proofs exist). Our only argument in favor of the existence of such zero-knowledge interactive proofs is based on much simpler assumptions, which are supported by strong intuitions (e.g., the existence of efficient processes that are hard to reverse). Indeed, at the last account, we do refer to intuition – but it is intuition regarding relatively simple and familiar phenomena (rather than intuition regarding complex and unfamiliar phenomena). The very fact that, at the current stage in history, we cannot claim a good understanding of the nature of efficient computation forces us not only to rely on assumptions but rather to differentiate between assumptions.

The foregoing example is far from being unique. The history of research in Cryptography is dominated by examples in which newly defined constructs were shown to exist based on simpler (and older) assumptions that enjoyed wide belief. In many of these cases, it was not a priori clear whether the newly defined constructs can at all be implemented, and the implementation based on better understood assumptions is still the only evidence to the feasibility of the relatively new constructs. We stress that in almost all of these cases, the assumptions being used enjoy also the belief of Koblitz & Menezes. It is something else that Koblitz & Menezes object to.

Let us discussed the objection of Koblitz & Menezes. Loosely speaking, they do not really object to any of the popular assumptions used in Cryptography. What they object to is the commitment of many researchers to the distinction between what they know for sure (based on rigorous analysis) and what they do not know (but may conjecture). In our opinion, maintaining this distinction should be commended (and represents a commitment to the methodology of science). In contrast,

---

intensive interaction.

[12] At a later stage of this essay, we shall argue that these definitions are sufficiently clear in the sense that it is evident that they address all reasonable security concerns. What we emphasize here that it is unclear *a priori* whether these definitions can be satisfied (let alone whether a specific system satisfies them). Furthermore, the feasibility of satisfying such definitions is significantly less obvious than the feasibility of satisfying simpler definitions such as of one-way functions.

the typical arguments of Koblitz & Menezes run as follows: *It is known that X holds* (possibly based on A that we all believe), *so why don't we just assume that X' holds too, because X' is closely related to X but is more appealing than X.* We note however that, unlike A, the assertion X is a fairly complex, and it is not clear what "closely related" means and whether this vague notion suffices to transport the truth value of X to X'. Let us be more specific.

Koblitz & Menezes consider several cryptographic schemes and refer to highly non-trivial security requirements such as semantic security under chosen-message-attack. In each case, X is an assertion about the security of some scheme, which is (rigorously) inferred based on some assumption A. The assertion X' refers to a modification of X, which Koblitz & Menezes consider insignificant (without providing any justification). In some cases, this modification is obtained by changing a single instruction in the algorithm (e.g., omitting one bit), but in other cases it is of fundamental nature (i.e., replaces an imaginary Random Oracle that lacks any structure by a specific "cryptographic hash" function such as MD5). Furthermore, Koblitz & Menezes (like anybody else at this age) have no real understanding why X hold, except that they know that X holds based on A. Nevertheless, they allow themselves to insist that X' must hold (provided X does).[13] They justify this bold statement by *their intuition* that the modification is insignificant, but they do not explain (let alone rigorously justify) why this modification is insignificant. In our opinion, without such an justification, the reasoning is bluntly flawed.[14] Furthermore, we stress again that their intuition refers to things (i.e., X) that they do not understand. Specifically, how can they insist that omitting a bit does not matter, when the currently known proof builds on this bit? How can they say that an unspecified "cryptographic hash" function is as good as an imaginary Random Oracle, while not being able to identify the structural property of a Random Oracle that is used in the known proof?

A side comment regarding the DDH assumption: Loosely speaking, the DDH assumption is an intractability assumption that refers to distinguishing between two types of probability distributions. While being empathic of Koblitz & Menezes's discomfort with the DDH assumption, we fundamentally differ regarding the source of discomfort. We are not concerned at all by the fact that, in "related" algebraic domains, assumptions regarding DDH fail. As hinted above, we view inferences based on superficial similarity (of the "related" type) as highly unsound.[15] What concerns us about the DDH assumption is the fact that this assumption refers to a setting that is less simple than usual (e.g., DDH is less simple than DH), which makes this assumption harder to evaluate.

[13] Needless to say, they do not really know if X' holds even provided that A does.

[14] We would not have protested against anybody claiming that X' holds. What we protest against is the argument that X' holds *because* X does. Ironically enough, Koblitz & Menezes wish to capitalize on the confidence attributed to X via the rigorous analysis methodology, and transport it to X', while basing the transportation solely on their intuition. Why don't they just offer the conjecture that X' holds *without referring to X at all*? Why do they invoke something that they object to (i.e., the proof that X holds)?

[15] A nice example, raised by Manuel Blum, refers to the fact that integer factorization is believed to be infeasible while the "related" problem of polynomial factorization is efficiently solvable. Another example refers to the relative complexities of computing the determinant versus computing the permanent, which are identical in the field GF(2) but are believed to be significantly different in the "related" field GF(3). Needless to say, 2SAT has linear-time algorithms, while the "related" 3SAT is NP-complete, etc, etc.

Rigorous analysis: the conclusion side

Having discussed the hypothesis side of the rigorous analysis of cryptography, we finally turn to its conclusion side. Indeed, a major oversight of Koblitz & Menezes is their failure to recognize that what distinguishes cryptographic research since the early 1980's from prior periods is not the use of rigorous proofs but rather what is being proved.

In a nutshell, the cryptographic research since the early 1980's proceeds in two stages: a definitional stage and a constructive stage. In the definitional stage the functionality underlying a natural security concern is identified, and an adequate cryptographic problem is defined. In such a definition, the desired functionality is defined in terms of operation in an imaginary ideal model, and a candidate system is required to emulate this operation in the real-life model, which in turn is clearly defined (while specifying the adversary's abilities). Only once the definitional stage is completed, one proceeds to construct a system that satisfies the definition. Thus, the rigorous analysis methodology is present in both stages.

As argued before, the definition obtained in the first stage is such that it is not obvious whether it can be met at all. Similarly, it is typically extremely hard to determine whether a given system meets this definition. Thus, one should infer the latter fact based on better understood assumptions. Furthermore, it is typically a good practice to design the system with an eye towards what is needed for establishing its security (i.e., a proof that the system meets the definition). Indeed, Koblitz & Menezes's objection to this practice is very odd, because even in the context of standard algorithms it is well-established that programs should be developed with an eye towards their proof of correctness...

Let us turn back to the definitional stage, and illustrate what may go wrong when it is left to the mercy of a non-rigorous state of mind. The following examples refer to possible definitions of secure encryption schemes, and are related to real stories. Clearly, it is not enough to require that, when applied with matching keys, the decryption operation is the inverse of the encryption operation. Nevertheless, it is often the case that people (implicitly) specify cryptographic schemes and evaluate them in analogous ways. Turning to more conscious attempts to define security, we first mention the requirement that it is infeasible to obtain the secret key (even when given many corresponding plaintext-ciphertext pairs). Note that this definition, which has appeared in some texts, says nothing about the security of the actual encrypted data (and is satisfied by the trivial "encryption" scheme that disregards the key and applies the identity transformation). Finally, we mention the requirement that it is hard to retrieve the plaintext from the corresponding ciphertext. This definition is also unsatisfactory, because it does not refer to the possibility of obtaining partial information about the plaintext (e.g., its first half, as in the case of an encryption scheme that only encrypts the second half of the plaintext and leaves the first part intact). Recall that, in contrast to the foregoing *ad hoc* attempts, a robust definition of security is obtained by comparison to an ideal model in a functionality providing perfect secrecy is postulated.[16]

The main reason that we listed all these well-known examples of unsatisfactory definitions is to stress the fact that, also in these cases, schemes were constructed and shown to satisfy the corresponding definitions. Thus, the point is not proofs (of "security") but rather what is being proved. Needless to say, the aforementioned schemes were later broken, but the amazing part of the story is that some people blame the proofs for this misfortune, rather than realizing that the

[16] In the context of private communication, one postulates the existence of a perfectly private channel that links the communicating parties and is inaccessible to the adversary. We admit that in the first published formulation of semantic security the use of the "ideal model paradigm" is only implicit, and refer the reader to the next section (which addresses the progress of science).

problem is one of inadequate definitions.

Another lesson that we wish to take from the foregoing examples is that "contrived" counterexamples (as we presented above) suffice for clarifying a conceptual problem. In fact, we argue that it is often the case that "contrived" counterexamples clarify the point better than realistic examples (which are likely to contain many details that tend to obscure the point). We believe that this phenomenon is illustrated by the "contrived" counterexamples presented above. Oddly enough, Koblitz & Menezes seem to view the presentation of "contrived" counterexamples as an indication to the lack of a "real" problem, and consider "contrived" counterexamples as an encouragement to hold the beliefs that the latter refute. Their reasoning and its flaws are briefly discussed next.

The reasoning of Koblitz & Menezes is that the presentation of "contrived" counterexamples indicates the failure of attempts to obtain non-contrived counterexamples. While this may be true in some cases, it may not be true in others (i.e., it may be the case that, for various reasons, the researchers did not seek non-contrived counterexamples).[17] More importantly, we find it strange to ignore the main message (i.e., that a counterexample, albeit contrived, exists) and be encouraged by the secondary message (i.e., that a non-contrived counterexample is not known yet). But most importantly, while the dichotomy of "contrived" versus "natural" is intuitive and appealing, we warn against basing a cryptographic scheme (or a cryptographic methodology) on it. In particular, this dichotomy is not robust (i.e., what is contrived in one setting may be natural in another), and lack of robustness is a big danger when adversaries are concerned. Indeed, concerns that were considered contrived at one time (e.g., the effect of totally unrestricted chosen-ciphertext-attack), turned out to harm the security of real-life systems.

The march of science: freedom, mistakes, and revisits

One feature of the scientific process was alluded to in the last paragraph. It is that researchers study what they choose to study rather than what other people would have wanted them to study (e.g., what the latter believe to be begging for study). This is one aspect of the so-called academic freedom. We stress that academic freedom refers to the choice of discipline and problems in it. In contrast, the method employed in the study is not free from the accepted methodology of the discipline (though, of course, one has the freedom to leave the discipline).

The march of science is thus the aggregate of the (relatively) free movement of many individuals. This reality implies that this march is not "linear" (i.e., it does not progress in one direction or at equal pace in all directions). Unfortunately, this march also includes mistakes (either by individuals or by the entire research community); this is regrettable but normal and unavoidable. Whoever does not like this fact should not take part in science (or in any other human activity, for that matter).[18] Needless to say, the occurrence of mistakes does not invalidate the scientific methodology but rather increases the importance of being committed to it; that is, the fact that a rigorous analysis may be flawed does not mean that one should abandon rigorous analysis but rather that one should apply it even more carefully.

Likewise, there is nothing wrong in re-visiting old problems (and old approaches) and discovering new perspectives on them. Such a re-visiting, even when it affirms views that were rejected before, is part of the progress. We note that the re-visited subject is always different from the way it was before (cf., "one cannot enter the same river twice" [Heraklitos]), because the field has evolved,

[17]The author of this essay may serve as an example, and hopefully this will not be considered contrived.

[18]Needless to say, mistakes occur much more frequently outside the domain of science. Specifically, intuition-based cryptographic schemes have failed way more often than schemes that were accompanied with clear definitions and rigorous analysis.

and in most cases the new perspectives could not have been reached (in the same form) during the previous visit.

On misconceptions or things becoming fetishes

As any other human activity, Science is not immune to misconception.[19] Typical misconceptions refer to a rigid interpretation of the insights of the scientific inquiry. One fundamental example is discussed next.

The theory of cryptography (as well as central parts of complexity theory) is commonly developed by referring to polynomial-time computations. While this is a very convenient convention (which is certainly indispensable when exploring new frontiers and wishing to abstract away as many details as possible), this convention is inessential to the theory of cryptography. (Thus, it is a conceptual mistake to identify the theory of cryptography (or complexity theory) with the "polynomial-time" convention.) Indeed, a more accurate (alas much more cumbersome) treatment of the same theory may refer to explicit resource bounds and provide a quantitative relation between resources that witness the related phenomena. Needless to say, such a treatment makes transparent the cost of the transformation claimed in a result (e.g., the relation between the security of a system and the hardness of a computational problem on which it is based). While the importance of the aforementioned cost to practice is commonly appreciated, its importance to theory is often disregarded (which is unfortunate because in many cases this cost reflects a phenomenon that is worthy of attention). Specifying the aforementioned cost also allows to characterize relations by their quantitative tightness (i.e., the cost incurred). Unfortunately, all the foregoing aspects are lost by those who wrongly believe[20] that the theory of cryptography only refers to polynomial-time computations and that it is oblivious of the replacement of one polynomial by another.

The foregoing example shows how a good idea (e.g., using a simplifying convention) may become a fetish, and by becoming a fetish cause some harm. A more specific example is the Random Oracle Model, originally suggested as a good *sanity check* but unfortunately misunderstood as a *yardstick* for security. Indeed, what happened with the Random Oracle Model reminds us of the biblical story of the Bronze Serpent, reproduced next.[21]

During the journey of the People of Israel in the dessert, the prophet-leader Moses was instructed by the Lord to make a "fiery serpent" as a symbolic mean for curing people that have been bitten by snakes (which were previously sent by the Lord as a punishment for some prior sin). Several hundred years later, the bronze serpent made by Moses has become an object of idol worship. This led the righteous King Hezekiah (son of Ahaz) to issue an order for breaking this bronze serpent to pieces. Let us stress that the king's order was to *destroy an object that was constructed by direct instruction of the Lord*, because this object has become a fetish. Furthermore, this object no longer served the purpose for which it was constructed.

This story illustrates the process by which a good thing may become a fetish, and what to do in such a case. Regarding the latter issue we emphasize two aspects that need to be evaluated regarding the situation at hand: the harm caused by the fetish versus the benefit that the object may still offer. We also emphasize that the action should be targeted directly against the fetish

[19] Indeed, the text of Koblitz & Menezes provides an example of several misconception regarding cryptographic research.

[20] Indeed, one may say in their defense that most texts contribute to this wrong impression by adopting the conventional focus on polynomial-time (and making no hint regarding the possibility of a more general treatment). Still, wise people should read beyond the literal meaning of texts.

[21] See *Numbers* (21:4-8) and *2 Kings* (18:4).

itself: In the story, the king destroyed the fetish itself, not the entire temple (where it stood). Likewise, when some aspect of the theory becomes a fetish, one should consider abandoning this aspect (but certainly not abandoning the entire theory).

So what about the two foregoing examples. Since the "polynomial-time" convention is very useful (and will certainly continue to be so) and since the damage caused by it is limited to a rigid interpretation of the theory by few people, we believe that issuing a warning about it is the right path of action. In contrast, in our opinion, the Random Oracle Model has caused more harm than good, because many people confuse it for the "real thing" (while it is merely an extremely idealized sanity check). Needless to say, as in the case of the bronze serpent, the blame is not with its creatures (who meant well), and the danger could not have been foreseen *a priori*. Still, given the sour state of affairs, it seems good to us to abolish the Random Oracle Model. At the very minimum, one should issue a fierce warning that *security in the Random Oracle Model does not provide any indication towards security in the standard model.*[22]

Lastly, on the role of intuition

While we have rejected intuition as a method for deriving inferences, we are far from underestimating the importance of intuition for the search of knowledge (i.e., Science). As a rule of thumb, intuition reigns where rigorous analysis cannot possibly enter. Note that we are referring to settings in which rigorous analysis is inapplicable in principle (i.e., cannot possibly be applied), and not to questions about which a rigorous analysis has (so far) failed to provide an answer. Needless to say, intuition has to be abandoned whenever a rigorous analysis shows that it is wrong. Let us elaborate.

There are two domains where rigorous analysis is clearly inapplicable. The first domain refers to some basic assumptions about the world and the possibility of reasoning about it (e.g., assuming causality and other "pure concepts of understanding"[23]), while the second domain is the "kingdom of free will" (e.g., determining our personal wishes). The first domain is clearly related to the founding questions of any discipline and its methodology. In particular, the discipline's methodology cannot assert its own validity, nor can it assert the importance of the discipline itself. The first domain also encompass the basic assumptions and models of the discipline (which are often coupled with its founding questions). For example, cryptography refers to some model of computation and assume some basic logic (i.e., rigorous inference rules). The suitability of these choices is a matter of intuition, and this is reflected by saying that the model is "reasonable" and the logical axioms are "natural" (or "self-evident").

The "domain of free will" plays a central role in evaluating the importance of a discipline to real life. For example, cryptography is important to the present society because people wish to maintain their privacy while other parties (viewed as adversaries) wish to violate it. Indeed, while these conflicting wishes may be explained by some other disciplines (e.g., sociology and/or psychology), they are assumed as intuitive axioms in cryptography (and cryptography does not investigate the reasons for these human conflicts (and cannot possibly do so)).

[22]Needless to say, nobody can deny that security in the Random Oracle Model *does not imply* security in the standard model. What we claim here is not merely the lack of implication in the relevant direction, but rather the lack of justification for claiming that security in the Random Oracle Model *per se* can be used as evidence to anything in the standard model. We also note that insecurity in the Random Oracle Model does not necessarily imply insecurity in the standard model, although this implication is valid (in the "natural" case) when the Random Oracle is replaced by a pseudorandom function.

[23]Indeed, see Immanuel Kant's *Critique of Pure Reason*.

We note that, similarly to the founding questions of a discipline, the most fundamental concepts of a discipline owe their central stature to intuition (and not to rigorous analysis). For example, we often say that computational indistinguishability (i.e., indistinguishability by any user or adversary) is of "natural appeal" (i.e., it appeals to our intuition). That is, Leibniz's postulate by which *indistinguishable things are identical* is only supported by our intuition.

Intuition may (and should) also guide our attempts to gain knowledge. N.B., it guides these attempts, but it cannot (and should not) replace them. That is, if we have an intuition regarding what should be true (or how can something be done) then we try to test and confirm this intuition, and are not content with the intuition itself.[24]

In contrast, although intuition plays a role in setting the inference rules, intuition cannot bridge a gap in the application of the inference rules. Whenever such a replacement is performed, the result is an *ad hoc* heuristic argument, which is clearly inferior to a rigorous argument. In such a case the question of justification arises, and relying on the same source of intuition for this justification is bluntly circular. Furthermore, *intuition invoked in time of need* (as any feeling that arises in time of distress) *is to be suspected, certainly not trusted.* For sure, such intuition cannot serve as a basis for knowledge, and things are even worse in cryptography (see our discussion regarding the danger in relying on intuition and *ad hoc* heuristics when adversaries are concerned).

[24]Needless to say, throughout the text we assume a fundamental difference between intuition (i.e., a speculation about knowledge) and knowledge (i.e., an apparent certainty about the subject).