On our duties as scientists (a personal version)*

Oded Goldreich†
Department of Computer Science and Applied Mathematics
Weizmann Institute of Science, Rehovot, ISRAEL
oded.goldreich@weizmann.ac.il

March 15, 2004

Abstract

Our primary duty as scientists is to contribute to the progress of science. By doing so, we best fulfill our duties to the society at large. Our primary duty is not fulfilled by merely doing research, but rather by communicating our findings to the relevant scientific community. The key role of this community in the scientific process, entails some important (secondary) duties towards this community (e.g., evaluation of scientific work and education of future generations of researchers).

Needless to say, this essay presents my own subjective opinions regarding the aforementioned issues. I also seize the opportunity to express some related and less related opinions as well as tell some personal stories.

*This version contains some personal stories, which some readers may find to be in poor taste. Reading this version is a commitment to pardon me for these stories. Alternatively, please refer to the main version, from which these stories are omitted.

†Written while the author was at Radcliffe Institute for Advanced Study at Harvard University, where he was partially supported by a Radcliffe Fellowship.

1 Introduction

There are no privileges without duties

Adv. Klara Goldreich-Ingwer (1912–2004)

It is often said that we, scientists, are a privileged "class", and indeed we are: We get nice salaries, we have relatively flexible duties, and we enjoy a fair amount of freedom. In fact, especially once we are tenured, we are almost unaccountable (e.g., we have no real boss or manager, and our *basic* salary does not depend on the quality of what we produce).

There is a good reason that this is the state of affairs. Skipping an adequate philosophical discussion, let me just say that it is hard to think of a good alternative: That is, I cannot think of an *effective* way of promoting science *without granting scientist the freedom to determine their own goals* (based on their understanding of their own discipline) and *entrusting the scientific community with the evaluation of scientific work*.

This state of affairs assumes that the scientists fulfill their duties, and in particular do their best in order to contribute to the progress of science. Indeed, failure of some scientists to fulfill this duty (and abuse the aforementioned trust and freedom) is often used as an excuse by people that are interested in changing the aforementioned state of affairs (and typically do not care too much about science). Thus, scientists that do not fulfill their duties are harming science not only by depriving it from their contributions but rather, by conspicuously abusing the system, they endanger the entire system.

Jumping ahead, I wish to clarify that inability to contribute in certain ways (e.g., do original research of significant importance) does not mean that one abuses the system. There are many ways to contribute to the progress of science, and each scientist should find the ways that best fit his/her abilities. When I talk of "abuse" I refer to people that fail to recognize their duty and/or to act in a way that is adequate in light of this duty.

2 The duties

I will claim that our primary duty as scientists is to directly contribute to the progress of science. By doing so, we best fulfill our duties to the society at large (see Sec. 2.3).

I will also claim that our primary duty is not fulfilled by merely doing research, but rather by communicating our findings to the relevant scientific community (see Sec. 2.1). The key role of this community in the scientific process, entails also some important (secondary) duties towards this community (see Sec. 2.2).

I wish to stress that the call of duty is fulfilled by a candid attempt to contribute as much as possible to the causes. It is not required that each scientist contribute to all causes, and it is not required that one contributes beyond one's ability. One is only required to search the ways in which one can contribute most and to contribute as much as one can (not more, but also not less).

I also wish to stress that the fact that some things are duties does not mean that one cannot enjoy doing them. On the contrary, a task (be it a duty or not) is best performed by a person that enjoys performing the task. Thus, if one does not enjoy fulfilling any of the duties required of a scientist, then one better seek a different profession.

2.1 Direct contribution to the progress of science

The primary duty of scientists is to contribute to the progress of science. This is fulfilled by studying questions that they consider of interest (i.e., "doing research"). The direction of research is determined by the individual scientists based on their own understanding of the current state of affairs in the relevant discipline(s). If they are lucky then they make interesting findings (i.e., "obtain results").

In contrast to what most inexperienced scientists think,[1] the story does not end in obtaining (interesting) results, but rather starts at this. The duty is then to communicate these results (or findings) to the relevant scientific community. Indeed, one cannot communicate interesting findings without obtaining them first, but other than that the real duty is to communicate newly acquired knowledge. That is, contributing to the progress of science requires communicating new knowledge and not merely obtaining it. In fact, not communicating new knowledge to the relevant scientific community is effectively equivalent to not obtaining it at all.

Of course, nobody is very strict about not communicating his/her results. People want to communicate their results, but they often do not understand that it is their duty to communicate their results clearly such that others may understand the results while investing significantly less effort than required for reconstructing these results from scratch. That is, it is the duty of scientists to provide *a clear and detailed exposition of their findings*. Let me end this subsection by reproducing some text from my essay on "How to write a paper":

> The purpose of writing scientific papers is to communicate an idea (or set of ideas) to people who have the ability to either carry the idea even further or make other good use of it. It is believed that the communication of good ideas is the medium through which science progresses. Of course, very rarely can one be sure that his/her idea is good and that this idea may (even only eventually) lead to progress. Still in many cases one has some reasons to believe that his/her idea may be of value. Thus, the first thing to do before starting to write a paper is to ask *what is the idea* (or ideas) that the paper is intended to communicate. An idea can be a new way of looking at objects (e.g., a "model"), a new way of manipulating objects (i.e., a "technique"), or new facts concerning objects (i.e., "results"). If no such idea can be identified one should reconsider whether to write the paper at all. For the rest of this article, we assume that the potential writer has identified an idea (or ideas) that he/she wishes to communicate to other people[2].
>
> Having identified the key ideas in his/her work, the writer should first realize that the purpose of his/her paper is to provide the best possible presentation of these ideas to the *relevant community*. Identifying the relevant community is the second major step to be taken before starting to write. We believe that the relevant community includes not only of the experts working in the area, but also their current and future graduate students as well as current and future researchers that do not have a direct access to one of the experts[3]. We believe that it is best to write the paper taking one of these less fortunate people as a model of the potential reader. Thus, the reader can be assumed to be intelligent and have basic background in the field, but not more. A good example to keep in mind is that of a good student at the beginning stages of graduate studies[4].

[1] Clearly, I cannot know what other people really think. But I can speculate on what they think based on what they say or based on how they behave.

[2] We leave the case of criminals that pollute the environment with papers in which even they can identify no ideas, to a different article...

[3] Indeed the chances that the experts (in the area) will be the ones that further develop or use the new ideas are the greatest. Yet, much progress is obtained by graduate students and/or researchers who became experts only after encountering these new ideas and further developing or using them.

[4] Ironically, the writers who tend to care the least about readers that are at this stage of their development (i.e.

Having identified the relevant community, we have to understand its needs. This community is undertaking the ambitious task of better understanding a fundamental aspect of life (in our case the notion of efficient computation). Achieving better understanding requires having relevant information and rearranging it in new ways. Much credit is justifiably given to the rearrangement of information (a process which requires "insight", "creativity" and sometimes even "ingenuity"). Yet, the evident importance of having access to relevant information is not always fully appreciated[5]. The task of gathering relevant information is being constantly frustrated by the disproportion between the flood of information and the little time available to sorting it out. Our conclusion is that it is the writer's duty to do his/her best to help the potential readers extract the relevant information from his/her paper. The writer should spend much time in writing the paper so that the potential readers can spend much less time in the process of extracting the information *relevant to them* out of the paper.

2.2 Contribution to the scientific community

The scientific community is the "carrier" of scientific progress (or the medium in which the latter takes place). Thus, the needs of this community impose duties on its members. These duties are arguably secondary to the primary goal of contributing to the scientific progress itself, but still they are important (as long as the scientific community relies on them). The most important of these duties are (1) the evaluation of scientific work, (2) the consolidation of the accumulated knowledge in easily accessible forms such as surveys and books, and (3) the education of future generations of researchers.

There is, indeed, a connection between the three aforementioned duties: The education of future generations of researchers relies on the consolidation of the accumulated knowledge, whereas the latter presupposes the evaluation of the accumulated knowledge (for verifying its correctness and determining its relative importance). We thus start by discussing the evaluation of scientific work.

Actually, given the economy of resources, scientific work is evaluated firstly for the purpose of allocating resources; that is, the community (through its agents) has to decide whether or not to include a certain work in the program of a conference and/or publish it in a journal.[6] The community (through its agents) also has to decide whether or not to grant individual scientists some resources (e.g., a certain position in a certain department or a certain research grant), and this decision reduces to the evaluation of scientific work (or, at earlier stages, more to the promise of such). As long as these decisions need to be taken, it is the duty of scientists to serve in bodies that make these decisions and/or help these bodies by relevant reviews.

We now turn to the consolidation of the accumulated knowledge. High level expositions (i.e., surveys and books) are a major tool that helps bridge the huge gap between the gigantic body of existing scientific knowledge and our limited time (which does not allow to even scan all the relevant literature). Thus, it is important that such expositions be written, but of course this does not require every scientist (or ever most scientists) to write such expositions. Expositions should

beginning of graduate school) are those who have just moved out of this stage. We urge these writers to try to imagine the difficulties they would have had if they had tried to read the paper, just being written, half a year ago...

[5] Of course, everyone understand that it is important for him to have access to relevant information, but very few people care enough about supplying the community with it. Namely, most people are willing to invest much more effort in obtaining a result than in communicating it. We believe that this tendency reflects a misunderstanding of the scientific process.

[6] Presentation (resp., publication) slots in conferences (resp., journals) are resources not only because these ventures are physically limited, but also (and more importantly) because the attention capacity of the scientific community is limited.

be written both for the internal use of the discipline and for the external use of other disciplines, but of course these two audiences are better served by different expositions.

Finally, we get to the education of future generations of researchers. In contrast to what some scientists seem to think,[7] graduate students and junior scientists are a liability not a resource. It may be (and indeed usually it is) a pleasant liability, but this does not change its nature as a liability. How to fulfill this educational duty is indeed a good question. My own opinion is that education is linked to the educator's true personality; that is, the learning experience is effected by the entire behavior of the educator and not only by the parts of his/her behavior proclaimed to be devoted to education.

2.3 Contribution to society at large

My claim is that scientists best fulfill their duties towards the society at large by fulfilling their primary and secondary duties (listed in Sections 2.1 and 2.2). Still there are additional duties that involves legitimate and reasonable needs of the society that only the scientists can fulfill, and thus are required to fulfill. These duties refer to the education of the society about science.

The most concrete and most dominant aspect of education is the education of practitioners in areas related to science. (This clearly holds for sciences that have an applied branch, but not only for them: The society needs people that are educated in variety of disciplines, ranging from scientific ones to arts and humanities.)

In addition, a vibrant society has to be enriched by new ideas. These new ideas may come from a variety of disciplines, ranging from the arts and humanities to scientific disciplines. (Indeed, here one typically reverses the order...) Put in other words, culture consists not only of the arts and humanities, but also of science. Here we refer to the intellectual contents of scientific disciplines and not to their technological impact (about which people are often more curious).

A big issue being ignored: I have avoided the famous discussion regarding the responsibility of scientists for the social impact of their work. This discussion typically arises with respect to *direct* technological applications of science (e.g., the Atom bomb), and thus seems less acute with respect to theoretical computer science. But my main reason for avoiding this discussion is different: In my opinion, this discussion refers to the duties of scientists as members of the society (i.e., as human beings) and not to their duties as scientists, and thus is not within the scope of this essay. That is, the issue is whether a certain duty belongs to the "ethical code" of a specific profession or to the general ethics of human behavior.

3 Non-duties

The current section is written out of personal annoyance at people who behave as if it is my duty to do certain things (which I do not view as falling within my duties). I'm not saying that one should not do the following non-duties; I'm just saying that doing them is not within the call of duty.

[7] Again, I cannot know what other people really think. But I can speculate on what they think based on what they say or based on how they behave.

3.1 Provide service to non-scientific bodies or causes

Scientists are often asked to serve on various non-scientific bodies that plan and/or oversee activities that are only remotely related to science. Furthermore, these bodies typically have underlying agendas that are not focused at science.[8] Examples of such non-scientific bodies include:

- Governmental advisory committees for the development of technologies. Indeed, technology is related to science, but the objectives and methods of technological development are very different from those of science. In addition, technological development requires different skills (which include but are not confined to knowledge of the relevant science). Skilled engineers (which are initially trained by scientists) rather than scientists are the relevant experts for such committees.

- Governmental advisory committees for (lower-level) education. Even when the topic includes scientific education, it requires expertise in lower-level education. The relevant experts are educators (with a good background in science if we are talking about scientific education), not scientists.

  Even participation in committees for higher education is not a duty in case there is a clear impression that the committee merely serves as a rubber-stamp (especially when it is to implement policies to which one does not agree).

- Advisory committees for various non-governmental (non-profit) organizations. These may include boards of non-scientific cultural ventures. It may be nice to have scientists on board, but it is not their duty to play this decorative role.

- Advisory committees for commercial firms.

In some cases there is a false pretense that the aforementioned committees are related to science, but actually the scientist is used merely as an educated person and/or as a source of respectability. (Of course, there are cases in which these committees really need a scientific advice, but these cases are very rare and I do not refer to them here.)

Scientists are often asked to provide service to scientific bodies that indulge in non-scientific activities. An excellent example is committees formed to award various prizes. I reproduce below my opinion on this activity:

> Typically, awards glorify and focus attention on one achievement or contribution or individual, while ignoring the wider context (which typically includes many contributions by many individuals). This approach may be rarely justified. (One may argue that singling out one contribution is justified in case of real scientific revolutions, but such revolutions occur very rarely and do not provide enough substance for a yearly award.) In contrast, in any given year, there are a few outstanding contributions, and (arbitrarily) selecting one of them for an award seems unjustified. Things are even worst, because the border between the outstanding works and the rest is not that clear. This is not merely saying that the work of the selection committee is very hard, but rather that the latter is faced with a task that makes no sense.
>
> Of course, one can claim the same holds with respect to hiring decisions and selection of a program to a conference. However, the difference is that in the latter cases (i.e., hiring and conference's program), resources are objectively limited and so a selection is unavoidable. Of course, if one institutes an award, then one creates a new limited resource (and a selection is

[8] Although it may be nice to contribute to society in this way, I claim that this is not a duty of a person in his/her capacity as a scientist.

again called for), but my point is that there is no objective justification for creating this award at the first place.

Typically, awards are actually a lottery (which is non-biased at best) among equals, and once the outcome is determine people treat the winners as if they were better. How does this process serve truth (or science)?

Awards are bound to be superficial and attract attention to the superficial. Needless to say, there is a superficial aspect in any human, the question is whether it should be encouraged. Indeed, I should confess that I have participated in such activities (i.e., made nominations and cared about their outcome), but I'm far from being proud about it...

3.2 Be subjected to tech-talk

Some people like to listen to many technical talks (be it during conference sessions or in casual discussions), and manage to process these huge amounts of information. Others, like myself, cannot handle too much contents, and prefer to select very carefully what they listen to. Both approaches are legitimate, but my guess is that it is either the case that people of the latter type are very rare or that their needs and preferences are usually ignored. I am talking about my personal experience in conference: I am often approached by people who for one reason or another want to describe some idea to me, and seem totally unaware of the possibility that my capacity of absorbing new ideas is limited. What is worse (and make this issue relevant to the current section) is that they sometimes behave as if they don't care whether I want to listen or not. That is, they behave as if it were my duty to listen (or at least this is the way I perceive their behavior).[9]

4 A personal perspective (or three stories)

An essay about morals calls for some self-examination (i.e., how do you stand with respect to your proclaimed standards). I am going to tell three stories about how I failed to fulfill my duties at real time, and how I tried to redeem the damage later.

4.1 The GMW87 paper

My worst crime is my failure to produce, in due time, a detailed version of the work "How to play any mental game or a completeness theorem for protocols with honest majority" [7]. The aforementioned extended abstract provides only a rough sketch of the ideas underlying the construction of such protocols, let alone a proof of their security.[10] This would be OK if that extended abstract would be coupled with an adequate (detailed) exposition made publicly available within reasonable time. Unfortunately, such an exposition (i.e., [3]) has only appeared a dozen years after the presentation of [7].

Things are made worse by the fact that the results in [7] were relied upon in many subsequent works, and that I got a lot of personal mileage out of them. Certainly, it was my duty to see that an adequate (detailed) exposition of these results be made publicly available within reasonable time. The fact that providing such an exposition was highly non-trivial (i.e., that writing such an

[9] It would have been much nicer if they *candidly* asked "would you like to hear an idea regarding ⟨topic⟩". In this dream-world, only if I say "yes" would they continue by a two/three-sentence description of the idea, and ask again whether or not I would like to hear more details, and so on.

[10] In fact, the paper even lacks a satisfactory definition of security.

exposition was a very demanding task) only makes things worse, because it indicates that a decent write-up was even more required in this case.[11]

The only thing I can say in my defense is that in 1998 (a dozen years after the presentation of [7]), I could not bear this situation any longer. Out of realization of the need for such an exposition and my duty to provide it, I undertook the task of writing [3].

I wish to stress that the issue at hand is not the fact that the paper has not appeared in journal form. At the current time (i.e., with web-posting being more accessible than journals and with the non-uniform quality of verification (via refereeing) provided by journals), I do not think that a journal publication is a "must" (although I do think that it is desirable). The point is providing a *publicly-accessible detailed exposition of the work*.

4.2 The EG83 paper

The work "On the security of multi-party ping-pong protocols" [2] was the center piece of my D.Sc. thesis. An extended abstract of it has appeared in *24th FOCS*, and a full version has appeared as a TR. This version was submitted to a journal (around the same time), "accepted pending revision" but withdrawn because we failed to produce an adequate revision in due time (because this was not high on our priority list, and rightfully so). So I find nothing bluntly wrong about this story, except that it feels very strange not to publish the center piece of one's Ph.D.[12]

À propos that work, it illustrates how my own views regarding writing have changed in two years (from 1983 to 1985). When I received the referee reports regarding the said submission, I realized how bad the original write-up was. I happen to have a record of my response to the editor, reproduced below:

```
Prof. M.J. Fischer,
Editor-in-Chief, J. ACM
  ...

                                                        July 3rd 1985
Dear Prof. Fischer,

Thank you for your letter dated April 13th 1985 concerning my paper
with Shimon Even ("On the Security of Multi-Party Ping-Pong Protocols").
I find the referees' comments very useful. I am ashamed to admit,
but the original manuscript is indeed poorly written.
I am currently preparing a new version,
but it may take several monthes before I finish.

                                        Sincerely Yours

                                        Oded Goldreich
```

4.3 Writing a book on the Foundations of Cryptography

Since the late 1980's, I had a clear feeling that a book regarding the foundations of cryptography (as developed in works like [8, 1, 11, 9, 10, 6, 12, 7]) is in great need. I started writing such a book

[11] That is, if providing an exposition was a very demanding task for me, then certainly figuring out the missing details was very hard for others.

[12] I guess that some people may think that publishing the center piece of one's Ph.D is a duty or a requirement. Anyhow, in view of the revived interest in this work, I've posted (on my web-page) fragments of the said revision (which was written in 1985) as well as a scanned version of the original TR.

in 1991, but only a decade later was the first part [4] published (whereas the second part [5] is being printed these very days). So, I have done my duty, albeit a bit late.

The decision to partition the work into two parts was instrumental towards its completion. In retrospect, I found a good rationale for this partition, but the original motivation was to finish first the easier-to-write part, which is less needed (because the difficulty of writing an exposition is a good indication to how much it is needed). Let me end this subsection by reproducing text from the preface to the second part [5]:

> Writing the first volume was fun. In comparison to the current volume, the definitions, constructions and proofs in the first volume are relatively simple and easy to write. Furthermore, in most cases, the presentation could safely follow existing texts. Consequently, the writing effort was confined to re-organizing the material, revising existing texts, and augmenting them by additional explanations and motivations.
>
> Things were quite different with respect to the current volume. Even the simplest notions defined in the current volume are more complex than most notions treated in the first volume (e.g., contrast secure encryption with one-way functions or secure protocols with zero-knowledge proofs). Consequently, the definitions are more complex, and many of the constructions and proofs are more complex. Furthermore, in most cases, the presentation could not follow existing texts. Indeed, most effort had to be (and was) devoted to the actual design of constructions and proofs, which were only inspired by existing texts.
>
> It seems that the fact that writing this volume required so much effort implies that this volume may be very valuable: Even experts may be happy to be spared the hardship of trying to understand this material based on the original research manuscripts.

References

[1] M. Blum and S. Micali. How to Generate Cryptographically Strong Sequences of Pseudo-Random Bits. *SIAM J. on Comp.*, Vol. 13, pages 850–864, 1984. Preliminary version in *23rd FOCS*, 1982.

[2] S. Even and O. Goldreich. On the Security of Multi-party Ping-Pong Protocols. Technical Report No. 285, Computer Science Department, Technion, Haifa, Israel, June 1983. Now available from `http://www.wisdom.weizmann.ac.il/~oded/eg83.html`. Extended abstract in *24th FOCS*, pages 34–39, 1983.

[3] O. Goldreich. *Secure Multi-Party Computation*. Unpublished manuscript, 1998. Available from `http://www.wisdom.weizmann.ac.il/~oded/foc.html` Superseded by [5].

[4] O. Goldreich. *Foundation of Cryptography – Basic Tools*. Cambridge University Press, 2001.

[5] O. Goldreich. *Foundation of Cryptography – Basic Applications*. Cambridge University Press, 2004.

[6] O. Goldreich, S. Micali and A. Wigderson. Proofs that Yield Nothing but their Validity or All Languages in NP Have Zero-Knowledge Proof Systems. *JACM*, Vol. 38, No. 1, pages 691–729, 1991. Preliminary version in *27th FOCS*, 1986.

[7] O. Goldreich, S. Micali and A. Wigderson. How to Play any Mental Game – A Completeness Theorem for Protocols with Honest Majority. In *19th STOC*, pages 218–229, 1987.

[8] S. Goldwasser and S. Micali. Probabilistic Encryption. *JCSS*, Vol. 28, No. 2, pages 270–299, 1984. Preliminary version in *14th STOC*, 1982.

[9] S. Goldwasser, S. Micali and C. Rackoff. The Knowledge Complexity of Interactive Proof Systems. *SIAM J. on Comp.*, Vol. 18, pages 186–208, 1989. Preliminary version in *17th STOC*, 1985.

[10] S. Goldwasser, S. Micali, and R.L. Rivest. A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks. *SIAM J. on Comp.*, April 1988, pages 281–308. Preliminary version in *25th FOCS*, 1984.

[11] A.C. Yao. Theory and Application of Trapdoor Functions. In *23rd FOCS*, pages 80–91, 1982.

[12] A.C. Yao. How to Generate and Exchange Secrets. In *27th FOCS*, pages 162–167, 1986.