
האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

ON NOT WANTING TO KNOW

by

EDNA ULLMANN-MARGALIT

Discussion Paper # 174

June 1998

מרכז לחקר הרציונליות
והחלטות האינטראקטיביות

CENTER FOR RATIONALITY
AND INTERACTIVE DECISION THEORY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel

PHONE: [972]-2-6584135 FAX: [972]-2-6513681

E-MAIL: ratio@sunset.ma.huji.ac.il

WWW: <http://www.ma.huji.ac.il/~ranb>

ON NOT WANTING TO KNOW

Edna Ullmann-Margalit *

Abstract

A common assumption of practical reasoning is that, in order to act rationally, agents are to act on the basis of the totality of evidence available to them. Common practice and introspection, however, suggest that people often do not want to know. The paper explores various aspects of the phenomenon of not wanting to know in an attempt to find out whether it is inherently unreasonable. The exploration leads, first, to weakening the principle of total evidence through replacing it with a rebuttable presumption in favor of additional knowledge. The sustainability of this presumption is then examined in light of the large variety of circumstances in which it seems to be reasonably rebutted. The alternative which in the end is recommended is to give up both the general principle and the presumption, and adopt instead something like a case by case cost-benefit approach, where the value of additional knowledge is matched up against its cost. In the process, the key notions of available knowledge, the value of knowledge, and the cost of knowledge are elucidated; also, separate attention is given to the question whether not wanting to know may sometimes be argued to be either morally required or morally reprehensible.

* Center for Rationality and Interactive Decision Theory, the Hebrew University of Jerusalem.

To appear in: Edna Ullmann-Margalit (ed.), Reasoning Practically, Oxford University Press, 1998 (forthcoming).

ON NOT WANTING TO KNOW

Edna Ullmann-Margalit

"He that judges without informing himself to the utmost that he is capable, cannot acquit himself of judging amiss".
John Locke, Essays Concerning Human Understanding,
book II, chapter XXI, sec. 67.

1. Introduction

Wanting to know seems natural, and in need of no justification. Wanting not to know seems less natural and in need of some justification. But a moment's reflection suffices to make us aware that there are many types of things we may not want to know, and it seems that justifiably so. We may wish to protect ourselves from cluttering our mind and our memory, or from boredom, or from pain; we may wish to preserve our faculties of creativity and of imagination; we may wish to avoid excessive cost or unnecessary involvements; we may wish to remain impartial, or to retain an element of surprise in our lives. All of these will on occasion make us not want to know certain things.

Thus you may not want to know the number of hair on your head, the telephone numbers of everyone in town, the exact details of the trip abroad your neighbor made or of the operation she has undergone. Perhaps you do not want to know certain things in advance, like what do all the critics think about the play you intend to see, or whether or not a surprise party for your birthday is being planned, or whether the baby you are about to have is a boy or a girl. Adopted children may wish not to know who their biological parents are. It is possible that you do not want to know precisely what your spouse is up to when you are away - and many of us may not want to know the details of how prisoners are treated in

Singapore.

The phenomenon of not wanting to know, then, seems secure: it exists all right. But we may want to ask ourselves whether and to what extent it can be defended. When is not-wanting-to-know, and when is it not, inherently unreasonable or irrational? What is the relation between not-wanting-to-know and self deception? Is there an issue of morality here?

Theories of rationality conceive of rational action as the employment of appropriate means for achieving a desired end. Rational action is thus thought of as a product of two vectors. One is the vector of belief, or knowledge, or probability. The other is the vector of desires, or wants, or utility. I act rationally when I act so as to promote what I want on the basis of what I know.

Elaborate theories are constructed in order to account for these two constitutive elements, of belief on the one hand and of desire on the other. One way of looking at the present enterprise is to see what happens when we drop the assumption that these two vectors in the parallelogram of action are independent and allow them to interfere with one another. There are at least two ways in which the vectors might interact. One would be captured by the question: do we (always) know what we want? The other would be captured by the question: do we - and should we - (always) want to know? It is the second of these two questions that I focus on here.

2. **The principle of total evidence**

Let us consider an important component of what I have referred to as the vector of belief. It is the principle, or the requirement, of total evidence: a rational person should believe the hypothesis supported by all available relevant evidence. Addressing this issue, Rudolf Carnap distinguished the logical question of how we are to determine the degree r to which evidence E supports hypothesis H , and the practical-methodological question of what we are to do given r (however it is determined). The first is a question of theoretical reasoning, the second is a question of practical reasoning. From the fact that hypothesis H is supported by evidence

E to the degree r nothing follows regarding action unless a further assumption is introduced, namely that E is the totality of evidence available to the agent. If E is indeed the totality of evidence available to the agent at time t , then, according to Carnap, the rational agent will accept the directive to believe hypothesis H at time t to the degree r and to act accordingly (e.g., to bet on H with odds not exceeding r). This Carnap calls the requirement of total evidence, tracing a distinguished historical pedigree for it, including Jac. Bernoulli (in his Ars Conjectandi (1713), as cited by Keynes, A Treatise on Probability (1973 [1921], p.345-6), and Peirce.

Now when Carnap requires of his rational agents to attend to the totality of evidence available to them, he means that it is irrational for them to ignore or to disregard relevant information which they have. (Information is relevant insofar as it affects the degree r to which E supports H.) He discusses three types of cases which are meant to convince us that this is indeed so. First we are invited to consider a judge who ignores information brought before him which is relevant to the attribution of guilt to the defendant. The second case concerns a businessman who, in evaluating a proposed deal, disregards information he has about some of the risks involved. And the third is a scientist who publishes results of experiments supporting his theory but neglects to publish - or to consider - results unfavorable to it. In all three cases, according to Carnap, the procedure adopted is not rational.

However, even a cursory reflection about these supposedly paradigmatic cases brings to the fore some problems which ought to lead to further reflection. Carnap may have been insufficiently familiar with the Anglo-American adversary legal system, in which there is of course a fundamental distinction between relevant evidence and admissible evidence. A judge (or jury) often not only may but must ignore information which is brought up during the trial and which is doubtlessly relevant but which is deemed inadmissible qua evidence on social-moral grounds (evidence based on hearsay, on a wife testifying against her husband, on illegal wiretapping, etc.). A similar point may be made with regard to Carnap's third case: in science too a distinction may have to be made between relevant and admissible data. Suppose scientists are presented with results from the twins experiments conducted by Mengele, the

monster doctor from Auschwitz. Should we consider them irrational - as distinct from immoral - were they to refuse to look at this information or to use it in their own research, even if it may be relevant to this research?

As for the businessman, we should note that there is much psychological literature which suggests that people as well as organizations "often expose themselves to risk because they misjudge the odds" (Kahneman and Lovallo 1993, p.24¹) - and they misjudge the odds, sometimes, because statistical knowledge "that is known to the forecaster will not necessarily be used, or even retrieved, when a forecast is made" (*ibid.*, p. 26).

This observation is discussed by Kahneman and Lovallo in connection with the distinction they make between forecasts which draw on the "inside" and those which draw on the "outside" view: an inside view anchors predictions on the specifics of the case at hand, on detailed plans, and typically on representative scenarios; the outside view ignores the details of the case at hand and is essentially statistical and comparative. For example, consider estimates of how long it would take me to finish an academic project (a paper, a book, the development of a new curriculum). My own inside view prediction is typically based on my carefully detailed work plan, and is typically overconfident: we generally tend to exaggerate our own control over events, and to underestimate the likelihood of obstacles coming our way, with the result that finishing our projects almost always takes us longer than we expected. An outside view prediction will be based on comparative statistics of how long similar projects - my own as well as other people's - took to complete.

The authors conclude that while the outside view is "much more likely to yield a realistic estimate, --- the inside view is overwhelmingly preferred in intuitive forecasting" (*ibid.*, pp. 25, 26). So in terms of our own concern, the point to be made is two-fold. On the descriptive level, we have the psychologists' findings that forecasters indeed ignore relevant (statistical) information available to them. On the normative level, the question may be raised whether this is always to be condemned as unreasonable. Is it not the case that the optimistic bias, which is based on ignoring unfavorable but realistic information and which leads to unrealistic bold forecasts, may yet be beneficial? Ample evidence suggests that optimism is in

fact instrumental as a causal factor which contributes to successful coping with challenges. Taking the broad view, considerations of "productive enthusiasm", entrepreneurial initiative, morale, persistence in the face of difficulty etc. often speak for the benefits of unrealistic optimism (Seligman, Learned Optimism, 1991).

Let us go back to Carnap's principle of total evidence and ask, what does it mean to say that evidence E is all the evidence available to agent A at time t?

Evidence 'available' to A naturally means evidence in A's possession, but it might also mean evidence accessible to A, evidence that A could acquire. Carnap concentrates on the first notion: he is concerned to establish the irrationality of ignoring information that one already has. My own focus, in contrast, is on the second notion: I am questioning the rationality (or otherwise) of acquiring information that one does not yet have. But before I move on to this second notion of available evidence, let me briefly note that even for Carnap's own purposes his notion of available evidence remains problematic. The admonition not to disregard information that is in my possession could be variously interpreted according to whether it is meant in the narrow sense of information that I am - or happen to be - aware of at t, or whether it is meant in some broader sense which might include, e.g., retrievable information that I might have, or in an even broader (and troubling) sense which includes the deductive closure of everything I know.

Consider now the idea that in order to act rationally one ought to act on the totality of evidence accessible to him or to her: "He that judges without informing himself to the utmost that he is capable, cannot acquit himself of judging amiss" (John Locke, Essays Concerning Human Understanding, book II, chapter XXI, sec. 67).

This idea too stands in need of further clarification. To get hold of all relevant information before I act makes sense, but has its price. The acquisition of further information is likely to be costly - in terms of monetary outlay, exertion, time spent, mental or emotional burden ("He that increaseth knowledge increaseth sorrow": Ecclesiastes 1, 18). So perhaps the principle of total evidence could now be recast so as to require of agents, if they are to act

rationally, to act on the basis of the totality of relevant information accessible to them at a price they consider acceptable, given their goals.

But here we are in danger of rendering the whole enterprise trivial. Since there will always be some price attached to the acquisition of further information, the mere fact that on some particular occasion you do not want to know something in itself attests to the fact that you judge the price of obtaining it or of knowing it unacceptable. Hence no instance of avoiding knowledge would be considered irrational, or unreasonable.

What this suggests is that something like a cost-benefit approach may be useful here, where the value of knowledge - both intrinsic value and instrumental value - is matched against its cost. The cost divides broadly into cost of acquisition and the cost of having, and each of these can be categorized in various ways. The cost of acquiring information includes, e.g., time, money, effort, boredom, unpopularity. The cost of having information is the cost of having to come to terms with it and live with it, which may include anxiety, agitation, shame, guilt, remorse, pain, injury to self-esteem, and more. Both kinds of cost may involve strategic losses as well (a point I shall return to later).

3. **The presumption in favor of additional knowledge**

In light of the discussion thus far, and bearing in mind that we are focusing on practical rather than theoretical reasoning, we may at this point consider weakening the principle of total evidence, and examining a presumption instead. The presumption to be examined establishes, for purposes of rational action, a generic bias in favor of acting on the basis of more knowledge rather than on less. To defend the adoption of the presumption in favor of being maximally informed amounts to defending the belief that following it will lead, in the long run, to results which are overall better, in terms of goal fulfillment, than the results of following its antithesis (i.e. a presumption establishing a generic bias in favor of acting on the basis of less knowledge rather than on more), or indeed better than the results of a case-by-case balancing (i.e. of following no rule or presumption at all).

An important feature of a presumption as distinct from a general principle or rule is

that it is rebuttable in concrete instances. In any concrete instance in which we consider applying the presumption there may be reasons, or counter-indications, which will caution us against applying it in that concrete instance. When this happens we say that the presumption is being rebutted in that particular case. A presumption is by its nature rebuttable: it has an implicit unless-and-until clause attached to it. (Thus, the presumption of legitimacy in the Anglo-American legal system establishes a generic bias in favor of treating the mother's husband as the father of her child, unless and until there are counter-indications in a specific case.) Irrespective of being rebutted, a presumption may be revised: the generic rule itself may come to be viewed as ill motivated, or as having outlived its usefulness, or as unfitting to changed circumstances, etc., and therefore may come to require revision. Revising a presumption may mean changing it, reversing it, or discarding it altogether.²

Can the presumption for additional knowledge be sustained? Is it reasonable to defend it as a presumptive principle of practical reasoning, or does it collapse under the weight of its counter-examples? The examples discussed so far suggest that in spite of the initial plausibility of the idea that acting on the basis of more knowledge better serves our interests in reaching our goals than acting on the basis of less, there is much that we do not want to know, and there are many situations where knowing less seems to serve our interests better than knowing more. Several categories of such cases emerged: cases involving useless or irrelevant knowledge; cases where the cost of acquisition is too high; cases where the emotional cost of having is too high; cases involving problems of impartiality; cases where knowledge incurs hedonic losses having to do e.g. with the loss of spontaneity or surprise. In addition, there are interesting cases where there are strategic problems from knowledge. Sometimes refraining from seeking knowledge may be strategically advantageous, or having more knowledge may be instrumentally dysfunctional in the sense that it may reduce the chances of success in achieving the desired goal³.

How are we to take all these cases where the presumption does not apply? Are we simply to view them as an accumulation of rebutting circumstances to the presumption? Or, alternatively, are we to say that the sheer volume of the types of circumstances where the

presumption does not apply, as well as their sometimes systematic nature, suggest that the presumption itself needs to be reconsidered? Since the adoption of its counter-presumption (i.e. a principled bias in favor of less rather than more knowledge) is plainly irrational, perhaps the presumption ought to be revised in the sense of being localized, that is, restricted to particular contexts. Or perhaps it is to be pronounced unsustainable altogether. The latter course amounts to recommending something like a case-by-case cost-benefit approach, where the value of additional knowledge is balanced ad hoc in every particular instance against the cost - in the broadest sense - of obtaining it (where 'obtaining' is meant to cover both acquiring and having).

Now this recommendation is no light matter. There is something odd about taking the bias in favor of more knowledge as if it were just a putative technical principle of practical reasoning. It is, after all, well entrenched in our culture. Ever since the Enlightenment, knowledge in our culture has been contrasted with articulated ideologies, captured as dogmas, on the one hand, and with unfounded popular beliefs, captured as superstitions, on the other. Our culture treats knowledge - as it does freedom - not only as intrinsically valuable but as incrementally valuable, i.e. that more of it is always better than less. To be sometimes in favor of knowing less - which means to be against always knowing more - may sound like being against motherhood or friendship: it may sound like going against the grain of the culture.

In considering the question of whether the presumption for more knowledge should be sustained as a principle of practical reasoning in spite of the large number and variety of cases where it seems not to apply, or whether to abandon it altogether, one needs to look at the justification of the presumption. Broadly speaking, a presumption may be justified in instrumental terms, or it may be justified in normative terms (or, sometimes, in both kinds of terms). The first has to do with the factual question of what, in the long run, works best. The second has to do with the evaluative question, on which side we had better "err". What we have seen in the case of the presumption in favor of additional knowledge is enough, I believe, to indicate to us that an instrumental justification is shaky, at best. One would be

hard pressed to show convincingly that if we adopt this presumption then, as rational actors, our interests would by and large and in the long run be better served than if we do not adopt it and proceed on a case-by-case basis instead. As for a normative justification, it seems to me that it could only be anchored in the kind of cultural considerations outlined in the previous paragraph - that is, the intrinsic as well as incremental value our culture assigns to knowledge, and the high-minded reluctance to see knowledge as just one factor among others which enter cost-benefit calculations. Yet, given that the context in which we consider this issue is far removed from the historical context of the cultural wars of the Enlightenment against the counter-Enlightenment, it seems appropriate to consider the question of the value of additional knowledge as a question which a rational actor, engaged in practical reasoning, faces on a retail and not on a wholesale basis. (Let me muse, as an aside, that our tendency to use the term 'information' rather than 'knowledge' in the context of practical reasoning and rational choice theory is not unrelated to the fact that, unlike 'knowledge', 'information' in our culture is not put on a pedestal, so to speak, and there is consequently less aversion to think of it instrumentally.)

In what follows I shall not proceed to amass direct arguments against upholding the presumption in favor of additional knowledge. I shall instead go on to explore further aspects of the phenomenon of not wanting to know, all aimed at enriching the texture of this phenomenon, and all designed to increase our skepticism about the sustainability of the presumption. In the process several more examples will be considered, and several distinctions will be introduced; the last section will take up the moral context.

4. **Control of knowledge: from 3rd to 1st person**

Let us consider the formula "I don't want to know x". As our first move, let us think of it as a special case of the more general formula "I don't want A to know x". This other-person formula states my wanting to prevent - or to protect - someone from knowing something; I shall be assuming here that what I do not want A to know is not some private information

relating to myself. Usually when this is the case, I may be said to display paternalistic attitudes toward that person, or to engage in a manipulative power relationship with that person. If the interests I wish to promote are the other person's then it's the former, and if the interests I wish to promote are my own then it's the latter. What changes when we now switch from the third-person, general, formula to the special, first-person case where the person A is myself - when I don't want me to know x, so to speak?

Well, to begin with the sophist-like query may be raised, How do I know what it is that I do not want to know unless I know it already? My response to this is that while there may indeed be cases where I do not know that there is an "it" I wish I had not known until it is too late and I already know it, there are other cases where a specific question is on the table and it is the answer to this specific question ("What have you done?", "What is my test result?") that I do not want to know. And it is these cases that concern us here.

Next, the matter may be raised of the possible difference between "I do not want to know x" and "I want not to know x". While such a distinction may certainly be drawn, and in some contexts possibly even to some advantage, for present purposes I take these two formulations to be interchangeable. Both my not wanting to know and my wanting not to know imply an active attitude on my part to avoid obtaining the knowledge in question. It is not indifference to knowing which is at issue here, but aversion to knowing.

Now when the person I want not to know x is myself, can I be said to want to prevent or to protect myself from knowing x? The drift of the discussion so far is that the answer to this question is yes, I can. But when this is the case we may not wish to retain the interpretation of the formula in paternalistic or manipulative terms. To the extent that we do retain the language of paternalism or manipulation even when we talk about my not wanting me to know, I suggest that we have self-deception cases in mind.

The phenomenon of self deception, while certainly pertaining to the field I am here concerned to chart, does not exhaust it. It is even possible to maintain that cases of self deception are, strictly speaking, not cases of not-wanting-to-know. They are, typically, cases where one does know something, and at the same time is concerned to conceal this

knowledge from oneself. (Or, one does know the general contours of something, and at the same time prevents oneself from being informed about its details.) This phenomenon merits, and has indeed received, special expert attention. I shall set it aside, and proceed to explore the remainder of the field.

Before moving on let us briefly consider the formula: "You don't want to know x". This is a second-person formula, where another person is advising you to shield yourself from knowing something. This formula helps underline a certain ambiguity of 'want', between a curiosity-sense of 'want' and an interest-sense of 'want'. You may not want to know x in the sense that you are indifferent to x, i.e. that you have no particular desire or curiosity to know x, as distinct from the case where you perceive it as being harmful to you, as going against your interests, to know x. Thus, I may suspect that knowing x would make you pointlessly envious or anxious or otherwise upset you, or that it would adversely interfere with your performance of the task you face, or that it may bias you one way or another where impartiality is called for. We may think of the blindfold Athena, the goddess of justice, as emblematic of this latter justification for not wanting to know.

Consider the following example. A defense attorney in a murder case may decide to avoid asking her client the direct question of whether or not he committed the crime. She will, instead, only ask him whether he wishes to plead guilty or not guilty. The attorney in such a case may say to herself "I don't want to know the answer to the question whether my client is guilty or innocent". She assesses that she will do a better and more professional job if she does not know whether or not she is defending an innocent person. Indeed, once she knows that the defendant did commit the crime, she is actually more restricted in her choice of defense strategies than she is when she does not know whether or not he committed the crime (regardless, of course, of what she happens to believe about him). Moreover, she may come to have ethical qualms about her job, which might further detract from her professional performance. So that knowing less here implies having a larger range of options to choose from - which is in normal circumstances taken to be the preferred situation for the rational

decision-maker.

Thomas Schelling (in The Strategy of Conflict) and Jon Elster (in Ulysses and the Sirens), among others, have demonstrated some of the ways in which "tying our hands" in advance of action, that is to say visibly blocking off some of the courses of action open to us in a decision situation, may actually further our interests and hence be rational. The attorney example suggests something similar with regard to knowledge: that there are occasions on which intentionally arranging for ourselves to have less rather than more knowledge is strategically advantageous for us.

5. **Control of self-regarding knowledge**

Next, consider the formula "I don't want to know x" when x is self-regarding knowledge.

Since the focus of the present discussion is on the question of justifying my not wanting to know, we are exempt from discussing such issues as the freedom of and equal access to information in general, or the right of the public to know. But an assertion of my right to know that which concerns me does fall within the purview of our discussion. Let me explain. I may wish to assert my right not to be denied access to knowledge that concerns me personally, to the extent that this knowledge exists and that somebody out there has access to it. To withhold from me information about myself, whether it regards my personal and family status, my health, my financial situation, my school-, army- or job records etc., is to impinge upon my autonomy. Whatever one wishes to read into the dictum that knowledge is power, it unequivocally asserts that to control knowledge is to engage power games. To withhold such knowledge from me is paternalistic at best and manipulative otherwise. (Is information about my children considered information about myself? This is a borderline case. I suppose that I have the right not to be denied access to information about my children until they reach a certain age. The other side of the same coin is that I may be denied access to certain information about myself as long as I am a child under a certain age.)

To defend one's right to knowledge concerning oneself is one thing, though, and to choose to exercise this right is quite another. The question we are now asking is whether or

not one has the privilege of not wanting to know such information, and what if any are the limits of this privilege. Is not an adopted child, who may have access to information about her biological parents when she reaches the age of 18, free not to acquire that knowledge? Does a patient, who is not denied access to the results of the medical tests performed on him, have to be informed about the results? These cases point to a conflict which may exist between autonomy and welfare, with respect to obtaining knowledge. While being autonomy-increasing, knowledge about oneself may at the same time be welfare-reducing

In many types of situations we do indeed wish to waive our right to self-regarding knowledge. In managing our lives, both externally and internally, there seems to me to be no prima facie reason why we should not have different tastes, which would express themselves in different preferences over our mental states in general, and over our states of knowledge in particular. People differ in the way they cope, say, with painful knowledge, and this will be reflected in their preferences over what they want or do not want to know about themselves.

The first headline on the front page of the New York Times on October 24, 1995 read: "If Tests Hint Alzheimer's, Should a Patient Be Told?" The story is about a gene, called apo E4, such that people with two copies of it have as high as a 90 percent chance of developing the disease by the age of 80. It is, the paper goes on to tell us, "information that nobody seems to know what to do with." The 51-year-old woman featured in the story is quoted as saying that in her family, two sisters and a brother went in for apo E tests after she learnt that she had two apo E4 genes, but four sisters refused, saying they did not want to know. She herself never hesitated about wanting to know about her apo E result; neither did her own doctor who "found the temptation to know irresistible."

This case raises fundamental practical and ethical questions for doctors and researchers whose perspective is that of the providers of information. As for the patients' perspective, which is that of the consumers of information (as it were), this case seems to me to raise no fundamental questions. It rather helps underline the point just made that, insofar as one is dealing with knowledge which pertains to one's own vital personal affairs but which one can do nothing about, people have different tastes and preferences as to whether or not they want

to possess it. What one person may find "irresistibly tempted" to know, another may be to afraid to know. And just as is the case with other tastes and preferences, there isn't very much more of an illuminative, normative nature to say here.

There would be something more to say about this case had we denied the assumption that there was nothing one could do with the apo E-related knowledge. Even if it is true, and remains true, that there is no way to prevent or treat Alzheimer's disease, one may well hold that there are ways to prepare for the near-certainty of developing the disease. There are ways in which one, and one's family, may plan for the future. Once this point is acknowledged, the case under consideration can no longer be taken as a case exhibiting mere differences of taste-like preferences among people regarding their own cognitive mental states. There could then be a cogent argument for the unreasonableness of the ostrich policy of not wanting to know.

In any event, we should be clear about distinguishing cases in which I give up being informed about myself where this is merely a reflection of my personal taste, from cases in which I forego being informed about myself where this is supposed to be instrumental to promoting my ends. The latter cases, but not the former, may be evaluated as reasonable or unreasonable. Such cases may include those where I give up collecting information because it is too costly to do so, others where I want to remain impartial, and still others where having the knowledge will trouble or pain me, or strategically hurt my interests.

For example, in a case like waiving my right to read recommendation letters written on my behalf as part of my promotion procedure, my very signing of the waiver clause may causally affect the nature and quality of the procedure in a way that is ultimately more favorable to me than had I insisted on access to my file. This is a neat case where I may want to know in the curiosity sense of 'want' but not in the interest sense of 'want'. When I participate in an experiment, it may be better for me, as well as for the experimenters, if I do not know whether or not I belong to the control group. A variant on this is a case where an educational experiment is conducted on my students, and I choose not to know too many details of the experimental design so as not to be biased in my own attitudes toward the students while the experiment lasts. Or think of Rawls's veil of ignorance in the original

position. One way to reconstruct this, I suppose, is to say that those entering the deliberation room are expected to ignore self-regarding knowledge that they have: they are to relinquish all knowledge about their own particular characteristics and position in society, so as eventually to come up with a more just - and presumably more rational - design for their societal institutions.

6. **Hindsight**

"Had I known it at the time, I would not have done what I did." This is a familiar enough phrase, with various substitutions for the "it". "Had I known the real nature of the tasks and responsibilities involved in being a chairperson of this organization (board, committee), I would never have agreed to become a candidate in the first place." "Had I known what enrolling in that particular army unit really means, I would never have signed the papers." And so on. Now it is possible that when you say "Had I known what it involved, I would never have done it" the implicature is "And I wish I had known it at the time". That is, you regret having done what you did: you acknowledge that you acted on the basis of insufficient knowledge, and as a result you judge your action to have been wrong, or suboptimal, in some sense. You reproach yourself for not having found out more, for not having been better informed.

But let us consider now the possibility that the implicature is "And boy am I glad I had not known it at the time!" That is, you acknowledge that you acted on the basis of insufficient knowledge, and, furthermore, you realize in hindsight that on the basis of fuller knowledge the course of action you chose was not the rational one to pursue and that you would therefore not have pursued it. Yet, you are happy with the result.

Of course, being a rational decision maker does not guarantee - and does not claim to guarantee - the best result on every occasion. And if you gamble against high odds without knowing them and luck out, your being happy with the result does not retrospectively make your gamble rational. But the interesting possibility to be pointed out here is the case where your not knowing contributed causally to the success of your action - whether or not it would

have been the rational course of action for you to take on the basis of full knowledge. Had you known, for example, how many candidates have failed the exam you are about to take, you might either have been discouraged from even attempting it, or lost the confidence necessary for passing it. Your very ignorance of the rate of failure was thus a factor that contributed to the good result. Doubt, uncertainty or ignorance may sometimes improve performance - as may unrealistic optimism, discussed above.

7. **Morality**

We come, finally, to consider the aspect of morality. There are in fact two distinct kinds of cases here, which work in opposite directions. One concerns the possibility that for reasons of morality you should forego some piece of knowledge, i.e. the possibility that wanting to know is morally reprehensible. The other concerns the possibility that for reasons of morality you may not exercise your privilege not to know some piece of knowledge, i.e. the possibility that not wanting to know is morally reprehensible. The existence of both kinds of cases should further strengthen us in our conclusion that the question of obtaining additional knowledge ought to be settled on the merit of each case and not on the strength of a general rule or presumption.

With regard to the first kind of cases, recall the notion of inadmissible evidence referred to earlier. Our system of law recognizes that sometimes evidence, while clearly relevant to the question of guilt, should be disallowed. Knowledge gained by torturing a witness is a case in point. It is the procedure by which the information is obtained which is objected to here, not its content (albeit that the content of confessions obtained under physical torture is often tainted anyway). This procedure is judged to involve unacceptable social and moral cost. Note that the notion of cost here is not construed in terms of the price paid to obtain the knowledge, but rather in terms of the social consequences that the use of such knowledge might have for the future - as well as in terms of the intrinsic moral cost of the violation of a basic human right that torture necessarily involves.

Part of the rationale for not admitting this sort of evidence is the causal influence this

ban is meant to have in discouraging torture in future cases. Similarly in the case of scientific research based on unethical experiments, for which Mengele's twin experiments provide an extreme example. The information such research contains may be accessible, and it may even be of scientific value. But since its very use may be construed to condone the methods by which it was obtained, it may be best, normatively as well as prudentially, to forbid its use altogether.

As for the second kind of cases, we note our familiarity with contexts where your saying "But I did not know" meets with the retort "Well, you very well ought to have known". Taken as a moral admonition, when is this retort justified?⁴ Are there cases where not wanting to know is morally reprehensible? In other words, are there case where there is a moral duty to know?

The case of Nazi Germany once again provides an extreme case in point. When Germans - especially those in positions of influence - say "But I did not know what was going on", we feel that this is no defense against charges of complicity. The general point seems to be that we have a prima facie duty to inform ourselves about what is being done by others in our own name. Insofar as one is a citizen of a democracy, one bears some sort of responsibility not just for one's own actions, but also for the actions done on one's behalf by one's representatives. How and to what extent one can be expected to discharge this responsibility is a different issue, and a very complex one. But the point remains that at the very minimum the strategic move of shutting one's eyes, of not wanting to know, may not be a morally permissible one⁵.

A related point has to do with one's obligation to those under one's care. I may shut my eyes and be negligent with regard to my own needs; I may not want to know the medical diagnosis of my condition and what the appropriate treatment should be, and - up to a certain degree of negligence - outside interference may well violate my autonomy. But when it comes to my responsibility toward my child or my aging parent who depend on me for their care and welfare, my not wanting to know the details of their medical condition and its treatment must be morally condemned. So once again this is a type of situation where the notion of

responsibility is bound up with that of knowledge in such a way that we have a prima facie duty to know.

A final observation, from a different angle: sometimes knowing something about somebody may put you under a special moral obligation to that person, even if this is not a person toward whom you have special obligations otherwise. The situation is perhaps analogous to the case when you happen to be passing by the beach where a drowning man calls for help: your mere presence there at that instant puts you under an obligation to do what you can to save him. Similarly, if a casual conversation with a stranger on a train turns into a confession where she starts telling you all, you may well feel that the more you know, the more freedom you are losing. You realize that you are about to lose your status of a stranger with respect to this person and that the bond that is being created between the two of you in virtue of this very conversation will make you morally obligated toward her. You may well feel, then, that as long as you don't know you don't owe. I therefore offer a distinction here, between knowing and being told. And my suggestion is that there are cases where the issue is my not wanting to be told rather than my not wanting to know. Moreover, some interesting cases involve neither my not wanting to know, nor just my not wanting to be told, but rather my not wanting to be told by x: being told by him (or her) may bind me morally against my will to a person I do not wish to be, and would not otherwise be, morally bound to. It may also establish a power relationship between us which is one I want to avoid.*

ENDNOTES

* I wish to thank Harry Frankfurt, Avishai Margalit, and Cass R. Sunstein, for valuable comments on previous drafts of this paper.

¹ Daniel Kahneman and Dan Lovallo, "Timid Choices and Bold Forecasts", Management Science Vol. 39 no. 1, 1993: 17-31.

² I have drawn here on Edna Ullmann-Margalit (1) "On Presumption", The Journal of Philosophy LXXX, no. 3, 1983: 143-164; (2) "Revision of Norms", Ethics 100, 1990: 756-767.

³ Known in game theory are cases, e.g., where a player who wants some information may nevertheless refrain from seeking it: he may not want it to be known that he wants to know - because this in itself may be compromising information about his own position.

⁴ In law it is often said that it is our business to know, or at any rate that ignorance of the law does not exempt from sanction. In fact, however, this situation is conceptualized within the law in terms of a fiction, or a presumption of knowledge, rather than in terms of a duty to know.

⁵ A delicate balance must in fact sometimes be struck here, between the duty to inform oneself on the one hand and moral urgency on the other. The duty to be well informed is sometimes used as a pretext for doing nothing, or at least for procrastination, when it is quite clear that atrocities are being committed. There are occasions in which, even if there is some unclarity about the precise details of the atrocities (and there usually is), to insist on total evidence before taking any action against them is a total abuse of the principle and must be morally objected to.