Appendices

Appendix 1: Theoretical Framework for Diagnosis

Scientific medicine, it may be recalled from Section 4.1.5, is characterized by a *rational theoretical framework* together with deployment, in such a framework, of substantive knowledge derived from science. The framework and knowledge pertain to diagnosis, etiognosis, and prognosis (Sections 5.2, 5.3, and 5.4). The theoretical framework of gnosis (dia-, etio-, pro-) should be rational even if the type of knowledge that its application calls for is not available. Illustrative of all of this is the theoretical framework of diagnosis.

The theoretical framework for diagnosis that still is commonly taken to be rational – by 'clinical epidemiologists,' most notably – was adduced over a half-century ago by a dentist together with a radiologist (Ledley & Lusted, *Science*, 1959). The basic idea was that the probability of the presence/absence of a particular type of illness is not knowable directly when known is the diagnostic profile of the case; that knowable, even in principle, only are the *reverse probabilities*, the probabilities of the manifestations profile of the case in the presence and absence of the illness in question; and that, therefore, the diagnostic probability at issue needs to be derived by means of *Bayes' theorem*: Let $\mathrm{Y} = 1$ and $\mathrm{Y} = 0$ represent the presence and absence of the illness in question, respectively; and let X represent the vector of statistical variates that were adopted to represent the manifestations-based diagnostic indicators, with $X = x$ this segment of the diagnostic profile of the case at issue. Bayes' theorem gives the 'posterior,' x-conditional probability of the presence of the illness, P'', as a function of the 'prior,' manifestationally unconditional counterpart of this, P', together with the likelihood ratio, LR, for the (vector-valued) datum, $X = x$:

$$\mathrm{P}''/(1-\mathrm{P}'') = [\mathrm{P}'/(1-\mathrm{P}')] \times \mathrm{LR};$$
$$\mathrm{LR} = \Pr(X = x \mid \mathrm{Y} = 1) / \Pr(X = x \mid \mathrm{Y} = 0).$$

O.S. Miettinen, *Medicine as a Scholarly Field: An Introduction*,
DOI 10.1007/978-3-319-19012-9

This theoretical framework for diagnosis is logically *untenable*. Bayes' theorem is correct, of course. But the LR is, in the theorem's application here, a malformed concept; and even if it weren't, diagnostic knowledge of that reverse-probability (or probability-density) form would represent generally unsurmountable epistemic challenges.

Any LR is the ratio of two conditional probabilities (or probability densities), and conditional probability is meaningfully defined only if the condition is singular in its meaning, in its bearing on the probability at issue, here $\Pr(X = x)$. Illness, when present at the time of diagnosis, is generally a highly nonsingular type of entity as a cause of its manifestations. And when the illness at issue is not present to explain the diagnostic profile, its alternative is by no means singular either. The LR in this formulation is, thus, generally devoid of express meaning, even in principle.

Let us nevertheless assume, counterfactually, that illnesses targeted for diagnosis generally are singular entities as explanations of the illness-manifestational profile, so that the LR in that Bayes' theorem formulation of diagnostic probability actually is a defined, single-valued quantity. Now the need is to appreciate that the various possible values of the LR are not subject to becoming known as to their magnitudes. Suffice it to think of $\Pr(X = x \mid \mathrm{Y} = 1)$. Valid study of this requires access to cases of $\mathrm{Y} = 1$ (i.e., of cases of the illness at issue) independently of the manifestations profiles $X = x$. This is, however, generally unimaginable. And if this weren't a problem, the enormous number of the possible profiles conditionally on $\mathrm{Y} = 1$ would generally make study of their respective probabilities wholly impracticable.

Despite its untenability, that theoretical framework for diagnosis, centering on Bayes' theorem, has been adopted with great enthusiasm by 'clinical epidemiologists.' Central in this naturally has been the particular topic of the updating of a '*pretest*' diagnostic odds, based on the pretest diagnostic probability, P', by a test result to the corresponding '*post-test*' diagnostic odds, involving the post-test probability, P'':

$$\mathrm{P}''/\left(1-\mathrm{P}''\right) = \left[\mathrm{P}'/\left(1-\mathrm{P}'\right)\right] \Pr(\mathrm{T} = \mathrm{t} \mid \mathrm{I}) / \Pr\left(\mathrm{T} = \mathrm{t} \mid \bar{\mathrm{I}}\right),$$

where $\mathrm{T} = \mathrm{t}$ denotes the result (t) of the test (T) at issue, and where I and $\bar{\mathrm{I}}$ denote the presence and absence, respectively, of the illness at issue.

Even though diagnostic tests generally produce quantitative results, in this particular context there is an eminent tradition to reduce the quantitative result to a binary one, either 'positive,' T+, or 'negative,' T-, which point, respectively, to the presence and absence of the illness in question. And a well-established terminology is in use to denote the probabilities of these binary results:

$$\begin{aligned} &\text{Sensitivity}: \ \Pr(\mathrm{T}+ \mid \mathrm{I}), \\ &\text{Specificity}: \ \Pr\left(\mathrm{T}- \mid \bar{\mathrm{I}}\right). \end{aligned}$$

In these terms,

$$\text{LR for T+} : \text{Se} / (1 - \text{Sp}),$$
$$\text{LR for T–} : (1 - \text{Se}) / \text{Sp}.$$

These ideas have been extended beyond actual test results in the diagnostic profile in the manifestational profile, and even to the risk indicators in the full profile, in terms of decomposition of the overall LR to the 'test'-specific LR_1, LR_2, etc.:

$$LR = LR_1 \times LR_2 \times LR_3 \times \ldots,$$

thus overcoming the serious epistemic problem alluded to above. This, in turn, has led to a *huge industry* of studies on the 'sensitivity' and 'specificity' of various 'tests,' and also of 'meta-analyses' of their results. In some of these original studies, the 'test' results are addressed in terms of more detailed, ordinal categories, and LR values specific to these are then addressed.

Textbooks of medicine, and practitioners similarly, have taken but token notice of these ideas and of the products of the industry based on these, and this is just as well. Linguistically, it makes little sense to say that taking the history of myocardial infarction is a 'sensitive test' for the presence of a fresh myocardial infarction, or that electrocardiography as to ST level is a 'specific test' for the presence of pneumonia, even though both of these ideas are consistent with that inclusive definition of diagnostic test and the two performance properties of it. And conceptually, the problems with 'sensitivity' and 'specificity' are even more serious than those with the Y-conditional probabilities of $X = x$ above. The added problem now is that they, and each LR_i based on them, should be conditional on the set of profile elements accounted for before the *i*th 'test.'

The genuine concept of *sensitivity* of a genuine diagnostic test would be the tendency of its result to change from what it would be (or was) in the absence of the illness to what it would be (or was) in the presence of the illness; that the test result tends to reflect the actual presence/absence of the illness at issue (as the test tends to 'sense' the presence of the illness) rather than, notably, reflect the risk to contract the illness. Thus, in the diagnosis about myocardial infarction, history of MI has no true sensitivity for a current, fresh MI, even though it is an important one among the diagnostic indicators, nor is the taking of this history a diagnostic test (contrary to the teachings of 'clinical epidemiologists'). And genetic testing (when the illness is not genetically defined) has no true sensitivity to the illness, however valuable the genetic risk information may be in the diagnosis about the illness.

A test's *specificity* to a particular illness cannot be given any tenable definition; but a particular *result* of the test can be specific to the *presence* of the illness, meaning that this result occurs only in the presence of the illness, that it is pathognomonic about the presence of the illness; and a result of a test can be specific to the *absence* of the illness, pathognomic about this. Examples of specific results of diagnostic tests include the specificity of a positive result of the troponin test

to the presence of (a fresh case of) myocardial infarction, and the specificity of a negative/normal result of the chest-radiographic test to the absence of pneumonia.

. . .

A *tenable* theoretical framework for diagnosis is implicit in the nature of the object of it. Diagnosis – a form of knowing – is pursued in a particular generic *domain* of this (e.g., in the domain of the adult with chest pain as the chief complaint; Section 5.2); and the object of diagnostic knowing – probabilistic – about the presence/absence of illness I is:

$$\begin{aligned}&\Pr(\text{I present, given the diagnostic profile of a case from the domain});\\ &\Pr(\text{I} \mid X = x),\end{aligned}$$

where $X = x$ denotes the *entire* diagnostic profile of the case.

In statistical terms, this is the probability, P, that a random variate, Y, indicating the presence/absence of I (Y = 1 if I present, 0 otherwise) takes on the realization Y = 1 conditionally on the realization $X = x$ of a set of variates based on the diagnostic indicators involved in the (entire) diagnostic profile of the case. And a common type of statistical model for this type of situation is the *logistic* one:

$$\text{Log}\,[\text{P}/\,(1 - P)] = B_0 + \Sigma_\text{i} B_\text{i} \text{X}_\text{i}.$$

In this theoretical framework, once the Xs are defined, the object of the requisite knowledge is the set of values of the parameters (B_0, etc.) involved in the adopted logistic model.

Different from the reverse probabilities involved in the theoretical framework of diagnosis centering on Bayes' theorem (above), this logistic framework is one of well-defined probabilities whose values are subject to study – by identification of a series of instances of the presentation domain, documenting Y and X, and fitting the model to the data. Diagnostic research in this rational theoretical framework is a prerequisite for scientific diagnosis; but it has barely begun, while research in the false framework of addressing diagnostic indicators' 'accuracy' – 'sensitivity' and 'specificity' – for particular illnesses continues to flourish (cf. above).

. . .

In the continuing absence of meaningful diagnostic research, diagnostic probabilities are set on the basis of *practitioners*' '*tacit knowledge*,' derived informally from personal experience with cases of various diagnostic profiles from particular domains of case presentation.

Formally, the idea in this is that once a diagnostician has had experience with cases from a given domain – cases of the two possible realizations of a Bernoulli-distributed Y in association with various realizations of vector X of profile variates – (s)he has thereby learned how to set $\Pr(\text{Y} = 1 \mid X = x)$ for a new case from that

domain of case presentation – informally, without the statistical framework sketched above and without, even, any thoughtful review of the cases.

This idea – that a diagnostician's personal experience with cases in which the truth about the presence/absence of a particular illness gets to be known teaches him/her how to set diagnostic probabilities for the presence of the illness – is by no means easy to justify on a-priori grounds (Section 8.1); and empirically, 'expert' diagnosticians' diagnostic probabilities are known to be quite discordant in the context of particular instances of diagnostic challenge. But the idea endures that personal experience with cases is the basis for diagnostic expertise – while textbooks of medicine don't teach the knowledge-base of diagnosis, and education in whichever discipline of medicine doesn't.

This idea of practice-based learning of diagnosis therefore should be subjected to *research in cognitive psychology*, which likely would show that the idea is a myth (Section 8.1), at great variance with the truths about the learning potential of the human mind in this particular, very challenging context (outlined above).

In this psychological research, the focus could be on a very simple logistic model, with only two diagnostic-indicator variates, one binary and the other quantitative. For select realizations of these variates (X_1 and X_2) in N instances ('cases'), the corresponding chance realizations of Y (0 or 1) would be produced under the adopted model. The N instances of realizations of the three variates (X_1, X_2, and Y) would be presented to the study subjects, for them to learn how $\Pr(Y = 1 \mid X_1 = x_1, X_2 = x_2)$ could be set in some new instances from the presentation domain, with known realizations of the Xs but the respective realizations of Y unknown. Then, some new instances of the case profile would be presented to the study subjects, and they would be asked to set the corresponding probabilities for $Y = 1$, based on the experience with N 'learning' cases. The variates and their realizations would be presented in realistic substantive terms, and the participants would be arranged to have a motive to learn how to set $\Pr(Y = 1 \mid X = x)$.

If the idea of diagnostic learning from personal experience would be shown to be untenable by such psychological research, this would provide a strong impetus for medical research on diagnostic probabilities – under well-designed logistic models. See Appendix 2.

The theory of diagnosis and diagnostic research I address more extensively in my most recent previous book (*Toward Scientific Medicine*. New York: Springer, 2014), along with the counterparts of these for etiognosis and prognosis.

Appendix 2: Researcher-Assisted Learning from Practice

Section 8.1 was an introduction to realism about practice experience as a basis of learning about diagnostic probabilities. The focus was on a very simple example (hypothetical) of experience relevant to this learning; and the point of it was illustration of a grim fact: informally experience-based 'tacit knowledge' about

Table A2 Experience in Table 8.1 of Section 8.1 replicated, with the same size

D_1	D_2	D_3	I	D_1	D_2	D_3	I	D_1	D_2	D_3	I	D_1	D_2	D_3	I
–	+	114	–	+	–	116	+	+	+	119	+	+	+	122	+
–	–	99	–	+	–	94	–	–	+	109	–	–	–	79	–
+	+	130	+	–	–	103	+	–	+	105	+	–	–	102	–
+	+	108	+	–	–	104	–	–	–	84	–	–	–	102	–
+	–	109	–	+	–	110	–	+	+	119	+	+	+	123	+
+	+	122	+	+	+	133	+	–	–	113	–	+	–	110	+
–	–	89	–	–	–	105	–	–	–	101	–	+	+	96	+
–	–	115	–	–	–	106	+	+	+	116	–	+	+	110	+
–	+	127	+	–	+	117	+	–	+	110	+	–	–	109	–
+	+	127	+	–	–	105	–	+	+	112	+	+	–	96	–

See text

diagnostic probabilities is prone to be very insecure, perhaps not even qualifying as knowledge (Section 4.5.1). And in Section 8.2 the related main point was that practice experience can be truly instructive to the practitioner only through the intermediary agency of the diagnostic researcher.

For understanding and acceptance of this principle of case-based ('problem-based') learning of diagnosis, to focus on this pivotal species of gnosis, the student of medicine may need to know, in broad outline, how the researcher thinks about the problem and how (s)he in this theoretical framework uses the data to derive estimates of the probabilities at issue.

The researcher, if suitably oriented, thinks about the diagnostic probability for the illness in question as a function of the set of diagnostic indicators in the domain of the case presentation (Appendix 1). Having designed the form of the function (s)he fits it to the data. And the thus-obtained empirical probability function (s)he provides to the practitioner concerned to learn from the cases (s)he knows and has documented for statistical 'analysis' (synthesis) – and others – for reading the profile-specific probability estimates as they are needed in practice.

As for the example in Section 8.1, a suitable design (below) of the probability function's form and its fitting to the data in the table (Table 8.1) there leads to objective estimates of the probabilities in the two cases specified on the bottom of the table. They are, respectively, 0.15 and 0.85.

But a note of *great caution* is in order about the probability function based on those data. For one, the data could just as well have been the ones in Table A2 here; and based on these data, the corresponding estimates of those two probabilities are quite different from the ones above: 0.07 and 0.92.

In comparing these two pairs of estimates it is important to appreciate that the two sets of data were for the first and the second set of 40 cases from the same domain, generated independently under a probability model that implied the *correct probabilities* for those two cases to be 0.25 and 0.83, respectively. The third set gave 0.20 and 0.19 (*sic*), while the fourth one gave 0.44 and 0.83; etc. – all of these results derived under the same model that implied the correct probabilities.

The lesson from this is very discouraging: the statistical estimates of the diagnostic probabilities, while objective, are very unstable – they have *very low reproducibility* – in repeat experiences; that is, for reliable estimates needed is very large experience. From the first 400 cases, the estimates for those two new cases became 0.26 and 0.81 respectively. (Cf. the correct values above.) And to be noted here is that the larger is the number of parameters in the model fitted to the data, the larger needs to be number of cases (data points) for any given level of reproducibility/precision of the result. In the example here, the model involved only five parameters, but the results from sets of 40 cases were very imprecise nevertheless.

The statistical model underlying, and fitted to, the data was quite typical for a situation like this: logistic for the probability that I is present, with the linear compound for the logit of this probability (of $Y = 1$) formulated as $B_0 + B_1X_1 + B_2X_2 + B_3X_3 + B_4X_2X_3$, where X_1 and X_2 are indicators of D_1+ and D_2+, respectively, and $X_3 = D_3 - 100$.

The data on D_1, D_2, and D_3 were generated under a stochastic model for their joint distribution. Pivotal in this was the distribution of X_2 (the indicator of D_2+) as Bernoulli (0.5). X_3 was Gaussian with mean $100 + 20X_2$, variance 100; and X_1 was equal to X_2 with probability to 0.75 (otherwise $X_1 = 1 - X_2$).

For the generation of data on the presence/absence of I, the parameters B_0 through B_4 in the logistic model for $\Pr(Y = 1)$ were –1.1, 0.0, 1.1, 0.05, and 0.03, respectively (with the Xs as defined above).

. . .

As at issue here is diagnostic *research*, some notes on the *validity* of this type of experience as a source of evidence about the magnitudes of the parameters (*B*s) in the statistical model are in order. Apart from the correctness of the data as such, the question here is about the validity of the set of cases of record; and in this, there are two issues: the validity of the subset of the known cases from the domain such that the truth about the I status got to be known (to some doctors), and the validity of the subset of these informative cases that got to be part of the experience – first- or second-hand – being used for learning by a particular doctor.

The learning set has not lost validity on account of the truth having become known with probability that depends on the diagnostic indicators that are being accounted for. But the set of cases that constitutes the basis of learning obviously must not be select on the basis of the known I status itself, nor on the basis of correlates of this not accounted for in the model.

There are subtleties regarding the validity of the database for diagnostic research beyond the outlines of this above, but these are not a concern in this introduction to medicine rather than to medical research. However, the student and the practitioner concerned to learn about diagnostic probabilities through experience in practice,

informally, without statistics, needs to understand that the issues of validity of the learning set of cases are involved in this informal framework as well.

Upon this brief visit to the domain of diagnostic research against the backdrop of what came up in Section 8.1 in regard to the 'tacit knowledge' about diagnostic probabilities gleaned from diagnostic experience, it should be clear that learning diagnosis – diagnostic probability-setting – via the intermediary of diagnostic research is possible, while learning it informally from practice experience is quite illusory. Thus, while patients need to seek doctors' learned help to understand the implications of such sickness as they are experiencing, doctors need diagnostic researchers' help to understand the implications of the cases they've encountered or are otherwise familiar with. But ideally they would have the requisite knowledge-base of diagnoses from the diagnostic species of quintessentially applied medical science (Section 1.4 and Appendix 3).

Appendix 3: Neo-Hippocratic Learning of Medicine

In ancient Greece, the art of Hippocratic medicine was taught – solo – by masters of it. And their students, in turn, felt an obligation to propagate the learning by teaching the art to students of their own (Section 6.4).

In a modern country, no one can be a master of 'the art of medicine,' qualified to teach all of modern medicine. 'The art of medicine' has become a hollow concept: that art no longer exists; it has evolved into various component arts/disciplines of medicine. (Cf. Preface and Section 1.5.)

The modern counterpart of the Hippocratic teaching of medicine in ancient Greece would be the teaching of a *particular, suitably-defined discipline* of medicine (Section 7.2) by a single master of this art or by a set of these (interchangeable) masters. Thus, a *singular type of teacher* would again be the source of all of the learning, though now specific to whichever particular discipline of modern medicine, following whatever propaedeutic studia generalia of medicine (Preface) and study of the concepts and principles specific to the discipline in question.

In this modern counterpart of the Hippocratic teaching of medicine there would be one *major novelty beyond the focus and content* of it. The Hippocratic art was originally propagated by oral teachings alone (the 'Hippocratic' corpus of writings developing only later). By contrast, the relevant educational content of (i.e., the knowledge-base of practice in) a modern discipline of medicine – focusing on, say, diagnoses in a particular segment of emergency medicine (cf. Section 7.2) – could be *codified in its entirety* and even made *accessible as needed* in the course of practice (through an 'expert system' imbedded in cyberspace; Section 5.6).

Hippocrates – this outstandingly great teacher – was a practicing doctor; and in the neo-Hippocratic framework of learning medicine, similarly, the discipline-specific teacher would be a practicing doctor in the discipline in question and, as

such, a modern 'role model' for the student about the practice proper – about the deployment of the discipline's 'expert system' etc.

The teaching-and-learning following all of the relevant preliminaries ideally would take place, ultimately, in the teacher's actual practice, even though it initially would be conducted in the medical equivalent of the flight-simulator setting of airline-pilots-in-training: As of the moment a case from the discipline's domain – possibly one with a particular species of gnosis as its sole concern (Section 7.2) – would come up, it would be classified according to the categories of presentation in terms of which the system is organized (e.g., a diagnostic system according to species of sickness; Section 5.2). Upon entry of the presentation category, the system would assume its role as 'the expert.'

The system would ask questions and the doctor would answer these, based on facts ascertained on the case. And based on these inputs together with the general medical knowledge codified in the system – in terms of gnostic probability functions (cf. Appendix 1) – 'the expert' would give the relevant gnostic probabilities to the doctor – to form the basis of the doctor's teaching the client about their health, in the prognostic context as an input to their decision about the choice of treatment (Section 5.7). Upon these gnoses, in the context of actual practice in particular, the master would teach the client on the basis of the gnoses – and teach the student about such teaching, on a case-by-case basis.

Given that the requisite knowledge for gnoses would be codified in the expert system, the student would not need to personally acquire and memorize that knowledge (which would be impossible). In consequence of this, as noted above, the discipline-specific learning would focus on substantive concepts and principles specific to the discipline at issue – perhaps also imbedded in the discipline's expert system (along with the facts-acquisition algorithms and the facts-conditional gnoses) – together with the *skills* needed in the ascertainment of the gnosis-relevant clinical (rather than laboratory-based) facts on the cases that may come up.

Regrettably, in my view, this has not been the result of the evolution of the teaching and learning of medicine from the time of Hippocrates to the present in this Information Age. In fact, the evolution has taken away from the focus on and development of the knowledge-base (gnostic) of medicine, as the teachers, and to an extent, through their influence, the students too, are preoccupied with and distracted by 'basic' medical research and its results in addition to and even instead of medicine itself (Section 2.4), while quintessentially applied, gnosis-oriented medical research has not been seriously cultivated (Section 1.4.2).

And so in closing here I leave to modern students of medicine one final proposition to weigh and consider (à la Francis Bacon; Section 2.2). It is predicated on the premise that quest for personal happiness animates all actions of humans (Sections 2.3 and 5.8.2), and that happiness in professional life flows from the virtue (Aristotelian) of excellence in it (Sections 6.1 and 6.3). My final proposition here is this:

A student aiming to achieve *excellence* in medicine could more realistically expect to find it – and its consequent *professional happiness* – in the direction of

the neo-Hippocratic learning and practicing I sketched above as a dream about the future, than in the framework of the prevailing realities of medical education (Section 2.4) and medicine proper. For, in medicine as it still is, a major detraction from professional happiness continues to be the proto-Hippocratic frustration with the art (per the aphorism in Section 1.1) – how long the study of it is as a proportion of the length of life post entry into the study; how treacherous experience in practice is as a teacher of the knowledge-base of the art (Section 8.1); and how difficult in the practice of the art is the use of judgements as a substitute for knowledge about the relevant truths (Sections 5.2, 5.3, and 5.4). But in the future I'm dreaming of (for time beyond mine), study of one's art of medicine would not be unduly long (due to focus on the relevant, as sketched above); treacherous experience of practice would not have to be relied on as the source of 'tacit knowledge' (cf. above); and ad-hoc judgements would be replaced by objective knowledge (accessed from 'expert systems'; cf. above) – so that excellence of the deployed knowledge-base would be the cornerstone of overall excellence in every doctor's professionalism (cf. Sections 9.1 and 9.2).

Appendix 4: Answers to the Questions in the Preface

Q 1: Are 'disease,' 'sickness,' and 'illness' synonyms?

A 1: These three terms are commonly used as though they were synonyms, but in critical usage they are not. 'Illness' is the antonym of 'health.' Sickness is what a person with a case of a hidden illness-definitional somatic anomaly may directly suffer from (or the unwellness may have an extrinsic direct cause). Disease, like defect and trauma, is a particular species of illness; it (different from defect) is a process-type somatic anomaly, and its pathogenesis (different from that of trauma) is intrinsic to (patho)biology. (Sections 3.2, 3.3, and 3.4).

Q 2: Are 'treatment,' 'therapy,' and 'intervention' synonyms?

A 2: These three terms, like those in Q 1, are commonly used as though they were synonyms, but in critical usage their denotations are distinct: Treatment (in medicine) is therapy only when intended to change the course of an existing case of illness (and not merely sickness from it) for the better. And treatment is intervention (therapeutic or preventive) only when it is an action directly on the client's soma (rather than a change in their environment or behavior). (Section 4.6).

Q 3: What is true and unique about all genuine disciplines of medicine (distinguishing them from all paramedical disciplines, i.a.)?

A 3: True and unique about – and hence definitional to – all genuine disciplines of medicine is the doctor's pursuit and (potential) attainment of first-hand esoteric knowing – gnosis (dia-, etio-, and/or prognosis) – about the health of the client, as the basis of all else in his/her services to the client (most notably teaching the

client about their health beyond the available facts pertaining to this, by bringing general medical knowledge to bear on those facts). (Section 4.1.1).

Q 4: Can a modality of treatment – surgery, say – be definitional to a genuine discipline of medicine?

A 4: Not justifiably. Disciplines of medicine are rationally defined according to what the gnoses in them are about; and under prognosis, all possible modalities are considered (Section 4.1.2). A discipline said to be one of 'surgery' is a discipline of medicine insofar as the relevant gnoses are in the essence of it and the use of surgical intervention by the doctor only an incidental element in it. (Sections 4.1.1, 4.1.2 and 9.1).

Q 5: What logically is the essence of diagnosis, and what is the source of its requisite knowledge-base?

A 5: Contrary to what is said in dictionaries of medicine, diagnosis cannot logically be said to be the determination of what illness is at the root of the patient's sickness. For one, it is not a generally-tenable premise that the client's complaint about sickness actually is a factual one and, as such, a manifestation of illness (instead of having an extrinsic direct cause). And even when it is, the set of available facts constituting the diagnostic profile as the ad-hoc input to diagnosis is not, generally, a sufficient basis for such a determination (i.e., pathognomonic about the nature of the illness underlying and causing the sickness). Logically, thus, the essence of diagnosis is knowing about the correct probability that a case of a particular hidden illness is present, given the diagnostic profile of the case. (Section 4.2.1).

The requisite knowledge-base of diagnosis consequently is, ultimately, about the probabilities of the presence of the various possible underlying illnesses, conditional on the diagnostic profile (in the form of diagnostic probability functions; Appendix 1). The source of 'tacit knowledge' about these probabilities is commonly thought to be the doctor's experience with cases presenting for diagnosis in his/her diagnostic practice. Belief in such practice-based learning about diagnostic probabilities is, however, quite unrealistic, the source of true diagnostic knowledge being diagnostic research. (Sections 1.4, 8.1 and 8.2).

Q 6: What are tenable conceptions of scientific medicine and medical science, respectively?

A 6: While there now are two eminent conceptions of scientific medicine (Section 4.1.5) and many others besides (Section 6.5), the only tenable conception of it is: medicine with a logical theoretical framework and, in such a framework, knowledge-base from (medical) science (Section 4.1.5).

Medical science is science aimed at advancement of medicine (its practice; Section 1.4). The intended advancement is either the availability of new 'tools' useful for medicine or improvement of the knowledge-base of medicine (Section 1.4).

Q 7: What is the essence of ethical medicine?

A 7: Medicine is ethical insofar as the doctor is worthy of the client's trust that the service provided by the doctor is, in every respect, excellent – that is, as good as anyone could provide. The essence of ethical medicine is the doctor's professional trustworthiness in this sense (Sections 6.3 and 7.3).

Q 8: How is professional happiness in medicine best assured?

A 8: This question is specifically about a physician's professional happiness as a doctor, and this needs to be thought of as his/her happiness about the quality – excellence (cf. A 7 above) – of his/her professional services to the clients.

Toward this happiness-producing virtue (Aristotelian; Section 6.1) needed first is commitment of excellence as a doctor, the first-order feature of this being uncompromised dedication to serving no ends other than the best interests of the clients (Sections 6.3 and 7.1). Given this commitment, the path to that goal of excellence begins with suitable definition of one's domain of the pursuit of excellence (Sections 1.5 and 7.2). And then, in the pursuit proper, the key is wisdom in the meaning of realism (Sections 1.1, 8.1, and 8.2). (In the future, I hope, all of this will be very different; Appendix 3 above).