Smart strategies for doctors and doctors-in-training: heuristics in medicine

Odette Wegwarth,[1,2] Wolfgang Gaissmaier[1,2] & Gerd Gigerenzer[1,2]

CONTEXT How do doctors make sound decisions when confronted with probabilistic data, time pressures and a heavy workload? One theory that has been embraced by many researchers is based on optimisation, which emphasises the need to integrate all information in order to arrive at sound decisions. This notion makes heuristics, which use less than complete information, appear as second-best strategies. In this article, we challenge this pessimistic view of heuristics.

METHODS We introduce two medical problems that involve decision making to the reader: one concerns coronary care issues and the other macrolide prescriptions. In both settings, decision-making tools grounded in the principles of optimisation and heuristics, respectively, have been developed to assist doctors in making decisions. We explain the structure of each of these tools and compare their performance in terms of their facilitation of correct predictions.

RESULTS For decisions concerning both the coronary care unit and the prescribing of macrolides, we demonstrate that sacrificing information does not necessarily imply a forfeiting of predictive accuracy, but can sometimes even lead to better decisions. Subsequently, we discuss common misconceptions about heuristics and explain when and why ignoring parts of the available information can lead to the making of more robust predictions.

CONCLUSIONS Heuristics are neither good nor bad *per se*, but, if applied in situations to which they have been adapted, can be helpful companions for doctors and doctors-in-training. This, however, requires that heuristics in medicine be openly discussed, criticised, refined and then taught to doctors-in-training rather than being simply dismissed as harmful or irrelevant. A more uniform use of explicit and accepted heuristics has the potential to reduce variations in diagnoses and to improve medical care for patients.

[1]Center of Adaptive Behavior and Cognition, Max Planck Institute for Human Development, Berlin, Germany
[2]Harding Center for Risk Literacy, Max Planck Institute for Human Development, Berlin, Germany

Correspondence: Odette Wegwarth, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. Tel: 00 49 30 82406 695; Fax: 00 49 30 82406 394; E-mail: wegwarth@mpib-berlin.mpg.de

INTRODUCTION

There was a time in history when diagnosing diseases was of little importance to doctors because virtually all patients, regardless of their illness, received the same treatments, such as blood-letting or cupping. Times, however, have changed. Today, making a diagnostic or treatment decision involves handling a large body of probabilistic information and processing it under pressures of time and a heavy workload. How do doctors manage this task?

Over the last decades, psychologists have examined how humans integrate probabilistic information into their reasoning under various conditions and how they should ideally do so. Much of the resulting work has embraced the idea of optimisation, which holds that all information available must be integrated in a defined manner in order for sound reasoning to take place; otherwise, second-best solutions are inevitable. One theory that has strengthened this belief and spawned many variants of replicating studies, both in the field of medical decision making[1,2] and elsewhere, is Tversky and Kahneman's[3-5] heuristics and biases programme. In psychology, heuristics are defined as simple decision-making strategies, also called 'rules of thumb', that make use of less than complete information. In order to conclude, however, that cognitive bias is at work when somebody uses a heuristic, one needs to set a prior norm of what constitutes sound reasoning. Within the heuristics and biases programme, this norm was defined by the laws of probability, and thus any deviation from these laws was defined as a bias. Although Kahneman and Tversky, who investigated the unconscious use of heuristics, initially considered that heuristics enable humans to arrive at mainly good decisions, they and other researchers advocating the heuristics and biases programme focused on the bias aspect only. This has led to the commonplace supposition that using less than complete information, regardless of whether this use is unconscious or deliberate, leads to non-optimal or faulty decision making. The medical community quickly adopted the heuristics and biases view[6-8] and left it largely unrevised until now. For instance, in the late 1990s, Elstein[9] still described heuristics as 'mental shortcuts commonly used in decision making that can lead to faulty reasoning or conclusions' (p 791) and blamed the practice for many errors in clinical reasoning. However, more and more researchers are beginning to realise, especially in fundamentally uncertain domains such as medicine, that expertise and good decision making involve the ignoring of some information.[10-14] But is the practice of ignoring information truly desirable in the context of making important medical decisions?

In this paper, we are going to challenge the negative view of heuristics held in both the psychological and medical communities. We focus on the deliberate use of heuristics in the design of tools that help doctors make good diagnostic and treatment decisions and demonstrate when and why using less than complete information represents a viable approach to medical decision making. We will end this article with a call for including the science of heuristics in medical education in order to curb the uneducated use of heuristics and thereby improve health care.

HOW SMART ARE SIMPLE HEURISTICS IN MEDICINE?

Diagnostic decisions: the coronary care unit

Imagine the following situation: a man is rushed to hospital with serious chest pain. The doctor suspects acute ischaemic heart disease and needs to make a quick decision: should the patient be assigned to the coronary care unit or to a regular nursing bed for monitoring? How do doctors make such decisions? And how should they?

One strategy is to rely on intuition. For instance, in a rural Michigan hospital, doctors sent some 90% of patients to the coronary care unit. Yet only 25% of patients admitted to the unit actually had myocardial infarction.[15] Similar results (ranging from 12% to 42%) were found in larger hospitals. This phenomenon is also known as 'defensive' decision making. It occurs in an environment where doctors can be sued for doing too little, but not for doing too much.

Given that defensive decision making leads to cost-intensive over-diagnosis and over-treatment, researchers at the University of Michigan Hospital tried to solve the coronary care unit problem by training the rural hospital's doctors to use a decision support tool based on logistic regression.[16] This tool, called the Heart Disease Predictive Instrument (HDPI), offers all relevant information in a combined and weighted form, yielding a chart with some 50 probabilities (Fig. 1).

If a doctor wanted to determine a patient's probability of having acute heart disease based on this chart, she needed to check the presence and absence of combinations of seven symptoms and insert the relevant probabilities into a pocket calculator. Yet although this procedure led to a systematic order of information through which it provided guidance,

Chest pain = chief complaint					
History	ST&T Ø	ST⇔	T⇧⇩	ST⇔	ST⇔&T⇧⇩	ST⇧⇩&T⇧⇩
No MI and no NTG	19%	35%	42%	54%	62%	78%
MI or NTG	27%	46%	53%	64%	73%	85%
MI and NTG	37%	58%	65%	75%	80%	90%

Chest pain, not chief complaint					
History	ST&T Ø	ST⇔	T⇧⇩	ST⇔	ST⇔&T⇧⇩	ST⇧⇩&T⇧⇩
No MI and no NTG	10%	21%	26%	36%	45%	64%
MI or NTG	16%	29%	36%	48%	56%	74%
MI and NTG	22%	40%	47%	59%	67%	82%

No chest pain					
History	ST&T Ø	ST⇔	T⇧⇩	ST⇔	ST⇔&T⇧⇩	ST⇧⇩&T⇧⇩
No MI and no NTG	4%	9%	12%	17%	23%	39%
MI or NTG	6%	14%	17%	25%	32%	51%
MI and NTG	10%	20%	25%	35%	43%	62%

Figure 1 The Heart Disease Predictive Instrument (HDPI), a decision-support tool in the form of a pocket-sized card (Source: [15]). ECG = electrocardiogram; ST = certain anomaly in electrocardiogram; MI = myocardial infarction; NTG = Nitroglycerin use for chest pain relief

many doctors disliked using the HDPI because of its complexity and lack of transparency.[17,18] What was the solution? Should these doctors have continued to classify patients according to (defensive) intuitions that were suboptimal but frugal, or should they have based their classifications on complex calculations that are alien but possibly more accurate?

Fast and frugal decision tree

Eventually, Green and Mehr[15] found an alternative to (defensive) intuition and complex tools: smart heuristics. These correspond to natural intuitions but can have the predictive accuracy of complex statistical models. An unexpected observation initially led hospital researchers to try a heuristic model. When studying the impact of the HDPI on doctors' decision making, the researchers noticed that once doctors had been introduced to the tool, which improved the quality of their decision making, its subsequent withdrawal did not affect the quality of their decisions: these, surprisingly, remained at the improved level. It was out of the question that the doctors might have memorised the probabilities on the chart or calculated the logistic regression in their heads. What else could have caused this effect? The researchers suspected that the doctors might, instead, have simply learned the important variables and that the quantitative computation itself was of little importance. This interpretation led to the deliberate construction of a simple decision-making heuristic for the coronary care unit allocation problem that used only minimal information and computation. Inspired by this idea, Green and Mehr[15] constructed a simple fast and frugal decision-making tree (Fig. 2). (For more details on the general properties of fast and frugal trees and
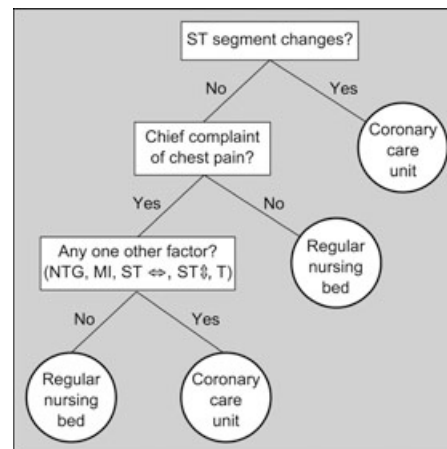

Figure 2 Fast and frugal decision tree for coronary care unit allocation (Source: [15]). ST = certain anomaly in electrocardiogram; MI = myocardial infarction; NTG = Nitroglycerin use for chest pain relief

their construction, see [19].) It ignores all 50 probabilities and asks only a few Yes/No questions. If a patient's electrocardiogram has a certain anomaly (the so-called ST segment change), he or she is immediately admitted to the coronary care unit. No other information is searched for. If that is not the case, a second variable is considered: does chest pain represent the patient's primary complaint? If not, the patient is immediately classified as low risk and assigned to a regular nursing bed. No further information is considered. If the answer is yes, a third and final question is asked to classify the patient.

How accurate is the fast and frugal tree?

Like the HDPI, the fast and frugal tree can be evaluated by multiple performance criteria. One of these is

accuracy, where the decision-making strategy should have, firstly, high sensitivity, so that it sends most patients who actually have a serious heart disease to the coronary care unit, and, secondly, high specificity, so that it sends few patients into the care unit unnecessarily. A second criterion is its ability to make decisions fast, which is essential when slow decision making can cost a life. A third criterion is frugality, which represents the ability to make good decisions with only limited information. The second and third criteria – speed and frugality – are inter-related and in both respects the fast and frugal tree is, by design, superior to the HDPI system, as might be doctors' intuition. So how accurate are decisions based on the fast and frugal decision-making tree compared with those based on the HDPI or on defensive intuition?

The answer is shown in Fig. 3. The y-axis represents the proportion of patients correctly assigned to the coronary care unit, as indicated by a subsequent heart attack; the x-axis represents the proportion of patients incorrectly assigned. The diagonal line represents chance performance. A point in the upper left corner would represent a perfect strategy, although that does not exist in the uncertain world of medical diagnosis. As the triangle shows, doctors' intuition initially performed at chance level or even slightly below it. The HDPI did better. Its performance is shown by squares, which represent various trade-offs between the two possible errors (false alarms, misses).

The fast and frugal tree, in turn, was more accurate than both doctors' intuitive judgement and the HDPI in classifying actual heart attack patients. It correctly


Figure 3 Accuracy of coronary care unit decisions made by doctors, according to the Heart Disease Predictive Instrument (■), (defensive) intuition (▼) and the fast and frugal tree (•). Accuracy is measured by the proportion of patients correctly assigned to the coronary care unit and the proportion of patients incorrectly sent to the unit. (Source: [15])

assigned the largest proportion of patients who subsequently had myocardial infarction to the coronary care unit. At the same time, it had a comparatively low false alarm rate. Note that the HDPI system used more information than the smart heuristic and could make use of sophisticated statistical calculations. Nevertheless, in this complex situation, using less information turned out to be of more benefit.

Treatment decisions: macrolide prescription

The heuristic approach has also been applied to target macrolide prescription in children with community-acquired pneumonia (CAP).[20] Macrolides represent the first-line antibiotic treatment for CAP, which is mainly caused by *Streptococcus pneumoniae*; infections caused by *Mycoplasma pneumoniae* are rare. However, macrolides no longer cover all bacterial causes of CAP. A study of schoolchildren in Pittsburgh found macrolide resistance in 48% of all group A streptococci isolated from throat cultures.[21] Given these alarming resistance patterns, the Active Bacterial Core Surveillance/Emerging Infections Program Network has urged doctors to reduce the inappropriate prescribing of macrolides, particularly to young children.[20]

Thus, after confirming a diagnosis of CAP in a child, the doctor must decide on the antibiotic prescription and further diagnostic testing. Although macrolides remain the antibiotic of choice in patients with *M. pneumoniae*, there are alternative antibiotics for other frequent bacterial infections. Rapid detection of *M. pneumoniae* is now possible by means of polymerase chain reaction analysis, but applying this test to all children with symptoms of CAP is costly. Moreover, most doctors prescribe a first-line antibiotic while they are awaiting the test result.

For such situations where time is crucial, information is uncertain and both costs and resistance rates need to be curbed, researchers[20] deliberately developed and tested two decision-support tools. One of these was a scoring system based on logistic regression. To ascertain a child's risk of having *M. pneumoniae*-triggered CAP with this scoring system, the doctor must verify the child's age and duration of fever, look up the respective scores for each of these in a table, and then sum up the scores before consulting an interpretation sheet. The other tool was a fast and frugal tree based on a heuristic approach and designed to help doctors rapidly identify the risk of *M. pneumoniae* as the cause of CAP in children. The fast and frugal tree (Fig. 4) adheres to the following heuristic rule: 'Prescribe macrolides only if the child is older than 3 years *and* has had fever for more than 2 days. Otherwise, do not prescribe macrolides.'[22]

How well did the two tools perform?

When doctors based their prescriptions on the scoring system, they were able to correctly identify 75% of all cases as being at high risk or very high risk for *M. pneumoniae.* The simple decision-making tree performed nearly as well: it correctly identified 72% of all cases as being at high risk or very high risk for *M. pneumoniae.* However, although both tools would help to curtail the superfluous prescription of macrolides to a considerable extent, the tree is yet more transparent: whereas the scoring system requires the user to look up data in a table, the fast and frugal decision tree, which asks, at most, two Yes/No questions, can easily be memorised.

MISCONCEPTIONS ABOUT HEURISTICS

These two examples reveal that common beliefs about heuristics are actually misconceptions. One of these misconceptions holds that humans use heuristics only because they have limited cognitive capacities. This often-repeated phrase incorrectly attributes the reasons for using heuristics exclusively to the limitations of the human mind, which is seen as an impoverished instrument. However, external reasons (e.g. that a problem is computationally intractable, the future is uncertain and the goals are ambiguous) can suffice for minds and computers to rely on heuristics. For instance, when former chess world champion Garry Kasparov played against the IBM supercomputer Deep Blue, both relied on heuristics, not only because both had limited capacities, but because the problem was computationally intractable: even the most brilliant

Figure 4 A fast and frugal tree for ruling out *Mycoplasma pneumoniae* infection in children with community-acquired pneumonia (CAP) (Source: [20]). AR = absolute risk; CI = confidence interval

minds and fastest machines were unable to compute its solution. Limitations of attention, memory and reasoning can, of course, contribute to the use of heuristics, but external reasons are sufficient.

Another misconception is that limited cognitive capacities are always bad. This belief is often implied but rarely stated, perhaps because it seems so obvious. However, although limited capacities may constrain functions, they may also, in fact, enable them.[23,24] For instance, large memory capacities in neural networks can prevent language acquisition in children, whereas starting small (limited capacity) and with simple sentences (baby talk) facilitates learning.[25] Luria's[26] famous mnemonist with almost unlimited memory could perfectly recall lengthy texts, but his memory was flooded by detail, making it difficult for him to summarise the gist of a text and think on an abstract level.

In comparison with optimising, heuristics are suspected of leading to second-best outcomes. If the optimal strategy is not known or too slow, however, using heuristics may well be the *only* solution. Moreover, every optimisation model is optimal only in relation to a set of mathematically convenient assumptions. Given that these assumptions do not hold in the real world, the outcome of optimisation can be disappointing; in such cases, optimisation theories are second-best.[11–13,27]

Another common misconception is that decision-making processes that use more information are always better than those that use less. In most models of rationality, it is taken for granted that the quality of decisions (or predictions) always improves – or at least cannot diminish – with an increasing amount of information. This assumption, however, is incorrect; the relationship between amount of information and quality of prediction is often illustrated by an inverse U-shaped curve.[28,29] Specifically, when uncertainty is high, as it is in numerous medical situations, the decision maker needs to ignore part of the available information in order to make robust predictions. For instance, in contexts where only a little information was available, the predictions made by a fast and frugal decision tree proved to be as robust as those supported by the benchmark of statistics, logistic regression, and only 1% point less so than decisions supported by the benchmark of machine learning, the classification and regression tree (CART), in various areas ranging from medicine to sports to economics.[19] Similarly, a simple strategy called 'take the best' was more accurate than complex strategies such as a CART and a neural network in making predictions in the majority of 20
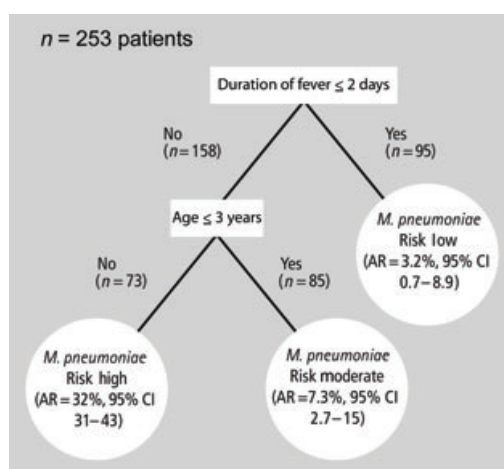
different decision-making situations.[30] Experts have been found to base their judgements on surprisingly little information,[31] and professional golf and handball players tend to make better decisions when they have less time to do so or when they act on the first idea that comes to mind.[32,33] But how exactly is this 'less-is-more' effect possible?

WHEN LESS IS MORE: ROBUSTNESS

To understand when and why less is more, it is important to understand the concept of robustness. In situations where decisions are liable to error − as they are in situations that involve uncertainty − robustness plays the key role in the less-is-more effect. The important distinction here is between data fitting and data prediction. *Data fitting* means fitting the parameters of a model to a body of data that is already known so that the model simply explains what has already happened. Here, using more information (free parameters) never hurts. By contrast, *data prediction* means testing whether a model can also predict comparable future events or outcomes. Here, however, using more information can hurt. If there are two diagnostic models, A and B, and A fits the known dataset better than B but predicts a comparable, yet new dataset less accurately than B, then model A is over-fitted to the known dataset. Over-fitting occurs when a model, by using too much information (free parameters), fits 'noise' and idiosyncrasies of the present dataset that do not generalise to a new sample. Yet, especially for situations whose structure is not known in advance, a model's most important feature is that it generalises well. A model's ability to predict (generalise to) new data – such as new patients – is called 'robustness'. Over-fitting, however, conflicts with the robustness of a model. To make the two concepts more transparent, suppose for a moment that you need a new dress. One means of meeting this need is to visit a tailor, who will take all your body measurements, assign these to the fabric you choose and create a dress that will fit you perfectly. That is what happens when a model is fitted to known data. Now suppose that a dear friend with similar general body features such as weight and size desperately needs a dress for an important event and asks if she can borrow yours. You, of course, agree. Your friend arrives at your door, eagerly tries on the dress, but sees that it does not fit her properly because some aspects of it are overly fitted to your body alone. This situation is akin to what happens when a statistical model is overly fitted to a specific set of data and is subsequently less able to predict another comparable set of data. By contrast, if you had chosen simply to buy an off-the-rack dress according to your size and weight, your friend might have been luckier: because of its less specific parameters, the dress would have been more likely to fit your friend as well. This analogy describes why a model that uses less information is more likely to generalise to comparable yet new data.

Like several other decision-related tasks in medicine, predicting heart attacks is far from error-free and no one case is 100% identical to another. In the original sample of several thousand New England patients on which it was validated,[16] the HDPI may well have provided a better fit than a fast and frugal tree. Yet, assuming that the predictive instrument is indeed an excellent tool for diagnosing patients in New England, it does not necessarily follow that it will perform equally well in Seattle, where new groups of patients will deviate in unknown ways from the original sample. In other words, the model that was best in the original population is not guaranteed to be best in these new populations. A fast and frugal heuristic that focuses only on the key variables is thus likely to be more robust and has a chance of performing better than the system that used more information. A world that is not perfectly predictable therefore requires that we ignore some information, as has been mathematically proven for specific situations.[30,34–36]

However, less information is not always better. Too little information can also be detrimental and eventually leads to under-fitting. In order to avoid both over- and under-fitting, a variety of methods have been developed to help us decide which of several models (e.g. decision-making support tools) has the right degree of complexity.[37] However, people seem to have a good sense of what information is important.[38] Although no general rule determines in advance how much and which information should be ignored, as a rule of thumb one can say that the more uncertain and the more redundant the information, the more of it should be ignored.[39]

THE (UNAPPRECIATED) POWER OF SIMPLICITY

Suppose that you regularly use the fast and frugal tree in Fig. 2 to allocate patients to either a care unit or a regular nursing bed. One of the patients you send to a nursing bed has a heart attack and dies. His relatives ask why the patient was not in the care unit and their lawyer finds out that you checked only two predictors and ignored all other information. The relatives sue you for malpractice. How many doctors are willing to take this risk?

The irony of the situation is that doctors often feel pressured to hide the ways by which they make decisions or to pretend the decisions were made on the basis of something more complicated. Part of this behaviour is rooted in the strong underlying belief that using heuristics will result in biases or in second-best solutions. The virtue of less-is-more is not yet fully understood and appreciated. As a consequence, the quality of treatment can suffer from covert and uneducated use of heuristics. In recent years, medical researchers have begun to see the potential of fast and frugal decision making and to appreciate it as a powerful alternative to the prescriptions of classical decision theory in patient care.[40]

However, any change in methodology must be supported by legal reforms that free doctors from the fear of being punished for doing the best they can for their patients. Effective litigation law would start from the simple insight that less *can* be more and that no medical decision is absolutely certain.[41]

Systematic training of doctors to use rules of thumb would allow them to make empirically sound, quick and transparent diagnostic decisions. McDonald[42] (p 56) emphasised this issue over a decade ago: 'The heuristics of medicine should be discussed, criticised, refined, and then taught. More uniform use of explicit and better heuristics could lead to less practice variation and more efficient medical care.'

Although we cannot present a complete curriculum describing how exactly the science of heuristics should be taught in medical education, what we can do is indicate some important milestones that should be met. Today's medical students should learn and understand that heuristics are neither good nor bad *per se*, but that their reliability and usefulness interplays with environmental circumstances, such as the inherent uncertainty of a specific situation. To broaden students' knowledge of what kind of environmental circumstances can be exploited in what fashion by what heuristic mechanisms seems as crucial as to teach them the building blocks from which heuristics can be constructed and adjusted for other problems or populations. After the basics have been delivered, a clinical teacher might continue, for instance, by introducing students to the various methods of constructing fast and frugal trees. In medicine, such trees are usually intended to first reduce misses and then decrease false alarms. This asymmetry will be reflected in the construction rules, which are aimed at achieving a large number of correct hits (e.g. correct assignments to coronary care units) at the first decisional level. For instance, one possible rule is to rank available information (e.g. chest pain) by sensitivity and to start the tree with the most sensitive piece of information. Practical units, where medical students can try out the success of different rules for self-chosen medical questions, will help to deepen students' understanding of heuristic tools and might even inspire novel research in the field of medical decision-making support tools.

As Green reported (personal conversation), doctors at the Michigan Hospital still enjoy using the fast and frugal tree, more than a decade after its use was initiated. Truly efficient health care requires that we master the complementary arts of focusing on what is important and discarding what can simply be ignored.

REFERENCES

1　Heller RF, Saltzstein HD, Caspe W. Heuristics in medical and non-medical decision making. *Q J Exp Psychol* 1992;**44A** (2):221–35.
2　Dawson NV, Arkes HR, Siciliano C, Blinkhorn R, Lakshmanan M, Petrelli M. Hindsight bias: an impediment to accurate probability estimation in clinicopathologic conferences. *Med Decis Making* 1988;**8**:259–64.
3　Tversky A, Kahneman D. Judgement under uncertainty: heuristics and biases. *Science* 1974;**185**:1124–31.
4　Kahneman D, Tversky A. Subjective probability: a judgement of representativeness. *Cognit Psychol* 1972;**3**:430–54.
5　Tversky A, Kahneman D. Availability: a heuristic for judging frequency and probability. *Cognit Psychol* 1973;**5**:207–32.
6　Bornstein BH, Emler AC. Rationality in medical decision making: a review of the literature on doctors' decision-making biases. *J Eval Clin Pract* 2001;**7** (2):97–107.
7　Klein JG. Five pitfalls in decisions about diagnosis and prescribing. *BMJ* 2005;**330**:781–3.
8　Detmer DE, Fryback DG, Gassner K. Heuristics and biases in medical decision making. *J Med Educ* 1978;**53**:682–3.
9　Elstein AS. Heuristics and biases: selected errors in clinical reasoning. *Acad Med* 1999;**74**:791–4.
10　DeMiguel V, Garlappi L, Uppal R. Optimal versus naive diversification: how inefficient is the 1/N portfolio

strategy? *Rev Financ Stud* 2009;in press, doi:10.1093/rfs/hhm1075.

11 Brighton H. Robust inference with simple cognitive models. In: Lebiere C, Wray R, eds. *Between a Rock and a Hard Place: Cognitive Science Principles meet AI-Hard Problems.* Association for the Advancement of Artificial Intelligence Spring Symposium; Technical Report SS-06-03. Menlo Park, CA: AAAI Press 2006;17–22.

12 Brighton H, Gigerenzer G. Bayesian brains and cognitive mechanisms: harmony or dissonance? In: Chater N, Oaksford M, eds. *The Probabilistic Mind: Prospects for Bayesian Cognitive Science.* New York, NY: Oxford University Press 2008;189–208.

13 Czerlinski J, Gigerenzer G, Goldstein DG. How good are simple heuristics? In: Gigerenzer G, Todd PM, ABC Research Group, eds. *Simple Heuristics that Make us Smart.* New York, NY: Oxford University Press 1999;97–118.

14 Reyna VF, Lloyd FJ. Physician decision making and cardiac risk: effects of knowledge, risk perception, risk tolerance, and fuzzy processing. *J Exp Psychol Appl* 2006;**12** (3):179–95.

15 Green LA, Mehr DR. What alters physicians' decisions to admit to the coronary care unit? *J Fam Pract* 1997;**45**:219–26.

16 Pozen MW, D'Agostino RB, Selker HP, Sytkowski PA, Hood WB. A predictive instrument to improve coronary-care-unit admission practices in acute ischaemic heart disease. *N Engl J Med* 1984;**310**:1273–8.

17 Corey GA, Merenstein JH. Applying the acute ischaemic heart disease predictive instrument. *J Fam Pract* 1987;**25**:127–33.

18 Pearson SD, Goldman L, Garcia TB, Cook EF, Lee TH. Physician response to a prediction rule for the triage of emergency department patients with chest pain. *J Gen Intern Med* 1994;**9**:241–7.

19 Martignon L, Katsikopoulos KV, Woike JK. Categorisation with limited resources: a family of simple heuristics. *J Math Psychol* 2008;**52**:352–61.

20 Fischer JE, Steiner F, Zucol F, Berger C, Martignon L, Bossart W, Altwegg M, Nadal D. Use of simple heuristics to target macrolide prescription in children with community-acquired pneumonia. *Arch Pediatr Adolesc Med* 2002;**156**:1005–8.

21 Martin JM, Green M, Barbadora KA, Wald ER. Erythromycin-resistant group A streptococci in schoolchildren in Pittsburgh. *N Engl J Med* 2002;**346**:1200–6.

22 Katsikopoulos KV, Pachur T, Machery E, Wallin A. From Meehl to fast and frugal heuristics (and back): new insights into how to bridge the clinical−actuarial divide. *Theory Psychol* 2008;**18** (4):443–64.

23 Hertwig R, Todd PM. More is not always better: the benefits of cognitive limits. In: Hardman D, Macchi L. *The Psychology of Reasoning and Decision Making: a Handbook.* Chichester: Wiley 2003:213–31.

24 Gaissmaier W, Schooler LJ, Mata R. An ecological perspective to cognitive limits: modelling environment−mind interactions with ACT-R. *Judgment Decis Making* 2008;**3**:278–91.

25 Elman JL. Learning and development in neural networks: the importance of starting small. *Cognition* 1993;**48**:71–99.

26 Luria AR. *The Mind of a Mnemonist.* Cambridge, MA: Harvard University Press 1968.

27 Bookstaber R, Langsam J. On the optimality of coarse behaviour rules. *J Theor Biol* 1985;**116**:161–93.

28 Pitt MA, Myung IJ, Zhang S. Toward a method of selecting among computational models of cognition. *Psychol Rev* 2002;**109**:472–91.

29 Gigerenzer G, Todd PM, ABC Research Group. *Simple Heuristics that Make us Smart.* New York, NY: Oxford University Press 1999.

30 Brighton H, Gigerenzer G. How heuristics exploit uncertainty. In: Todd PM, Gigerenzer G, ABC Research Group, eds. *Ecological Rationality: Intelligence in the World.* New York, NY: Oxford University Press. In press.

31 Shanteau J. How much information does an expert use? Is it relevant? *Acta Psychol* 1992;**81**:75–86.

32 Beilock SL, Carr TH, MacMahon C, Starkes JL. When paying attention becomes counterproductive: impact of divided versus skill-focused attention on novice and experienced performance of sensorimotor skills. *J Exp Psychol Appl* 2002;**8**:6–16.

33 Johnson JG, Raab M. Take the first: option generation and resulting choices. *Organ Behav Hum Decis Process* 2003;**91**:215–29.

34 Akaike H. Information theory and an extension of the maximum likelihood principle. In: Petrov BN, Csáki F, eds. *Second International Symposium on Information Theory.* Budapest: Akedemiai Kiado 1973;267–81.

35 Foster M, Sober E. How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions. *Br J Philos Sci* 1994;**45**:1–35.

36 Geman SE, Bienenstock E, Doursat R. Neural networks and the bias/variance dilemma. *Neural Comput* 1992;**4**:1–58.

37 Marewski JN, Olsson H. Beyond the null ritual: formal modelling of psychological processes. *J Psychol* 2009;**217**:49–60.

38 Pachur T, Bröder A, Marewski J. The recognition heuristic in memory-based inference: is recognition a non-compensatory cue? *J Behav Decis Making* 2008;**21**:183–210.

39 Dieckmann A, Rieskamp J. The influence of information redundancy on probabilistic inferences. *Mem Cognit* 2007;**35**:1801–13.

40 Elwyn GJ, Edwards A, Eccles M, Rovner D. Decision analysis in patient care. *Lancet* 2001;**358**:571–4.

41 Engel C, Gigerenzer G. Law and heuristics: an interdisciplinary venture. In: Gigerenzer G, Engel C, eds. *Heuristics and the Law.* Cambridge, MA: MIT Press 2006;1–16.

42 McDonald C. Medical heuristics: the silent adjudicators of clinical practice. *Ann Intern Med* 1996;**124**:56–62.