Title: A PROBLEM FOR ACHIEVING INFORMED CHOICE[‡]

Author: Adam La Caze

Philosophy Department,

University of Sydney

Abstract

Most agree that, if all else is equal, patients should be provided with enough information about proposed medical therapies to allow them to make an informed decision about what, if anything, they wish to receive. This is the principle of informed choice. It is closely related to the notion of informed consent.

Contemporary clinical trials are analysed according to classical statistics. This paper puts forward the argument that classical statistics does not provide the right sort of information for informing choice. The notion of probability used by classical statistics is complex and difficult to communicate. Therapeutic decisions are best informed by statistical approaches that assign probabilities to hypotheses about the benefits and harms of therapies. Bayesian approaches to statistical inference provide such probabilities.

[‡] This is a preprint of La Caze A. (2008). A problem for achieving informed choice. *Theoretical Medicine and Bioethics*, 29:255–265. The final version is available at http://dx.doi.org/10.1007/s11017-008-9069-x.

If all else is equal, you should be provided with enough information about proposed medical therapies to allow you to make an informed decision about what, if anything, you wish to receive. This is the principle of informed choice. Importantly, it is you who needs to choose. Informed choice and shared models of clinical decision-making can be contrasted with paternalistic approaches that hold the clinician as sole, or at least principle, decision maker. There is plenty of room for debate within the notion of informed choice.[1] Much less contentious is the idea that underpins informed choice, and the closely related, informed consent: it is the *patient* who needs to understand, and accept, any proposed therapy. The amount and type of information required to achieve this acceptance will vary from person to person. But central to all discussions is information on the likely benefits and possible harms of proposed therapies.

This paper raises a problem for achieving informed therapeutic choice. I argue that, even in ideal situations, adequately informed choice is substantially impeded by the way benefits and harms of therapies are analysed and presented.

To make an informed choice, I suggest, you need information on the probability of benefit and harm of the therapy. In the case of therapeutic interventions, this information is usually drawn from randomised controlled trials. Such trials are analysed according to classical statistics. This raises a problem: classical statistics does not provide the probabilities needed to inform choice. The conceptual baggage that accompanies the observed results of trials is typically ignored or misinterpreted. This problem is distinct

from, but plausibly exacerbates, the well-documented difficulties people have in correctly understanding information presented in probabilistic or statistical form.[2]

My intention is to outline the problem classical statistical analysis poses for achieving informed choice. This requires discussion of three intersecting fields of inquiry that are typically considered separately: the ethics of informed choice; the communication of risk; and the philosophy of statistics. First, I briefly review the notion of informed choice, and arguments in its favour. I argue probabilistic information on the harms and benefits of interventions is needed to adequately inform choice. A brief discussion of Bayesian statistics is provided to outline an approach that provides such probabilities. Second, I illustrate contemporary approaches to risk communication in health. In particular, I show that quantitative communication of risk relies on estimates provided by classical statistics. In the final section, I examine the notion of probability employed by classical statistics to show why it is far from ideal for informing therapeutic choice.

THE ETHICS OF INFORMED CHOICE

Informed choice is closely related to informed consent. The importance of informed consent, particularly in the context of participation in medical research, is well recognised, both morally and legally. Informed consent is a central component of the Declaration of Helsinki.[3] It ensures all candidates for medical research are given the free choice to accept or decline participation, at any stage of their involvement. To permit this free choice individuals need to be adequately informed of, among other things, the

potential benefits and risks of involvement. Informed choice extends these notions to routine clinical encounters.

Robert Veatch stresses the importance of the distinction.[4] Whereas informed consent alludes to information provided to an individual to allow them to opt in or out of a study (or for or against taking a particular treatment), informed choice alludes to the information needed for an individual to be able to appropriately choose between therapeutic alternatives within a clinical encounter. Veatch is concerned that "informed consent" implies that the clinician judges the "treatment of choice" from among the alternatives and then informs the individual of this choice. This would deny an individual the opportunity to consider viable treatment options on their terms. It also presumes that the clinician alone is in the best position to weigh the relative value of treatments for an individual. In some situations this may be the case; in most it is not. For this reason Veatch recommends abandoning "informed consent" in the clinical context.

It does not seem necessary to construe "informed consent" within clinical encounters as excluding "informed choice" (as Veatch does). However, I will use the distinction to emphasise the move from a research to a clinical context. Here "informed choice" will refer to the moral obligation of a clinician to provide sufficient information to permit an individual to choose from among the viable therapeutic alternatives (including no therapy). I take this obligation to be uncontroversial—at least, when the individual and

well enough, and "rational" enough (in a minimal sense), to be able to engage in the discussion.

The obligation to provide sufficient information to allow for informed therapeutic choice is typically voiced in terms of respect for an individual's autonomy. For example, informed choice is seen as, "the right of individuals to exercise control over aspects of their lives that they deem critical for whatever reason".[5] Perhaps even more directly, the clinician's obligation to inform choice can be grounded in an obligation not to deceive or coerce.[6] Omitting to provide enough information to a capable individual to make an informed choice amounts to a form of deception or coercion, and therefore cannot be justified.

While the moral imperative for informed choice or consent is clear, achieving it is difficult. Mostly, this is because we are not ideal rational agents. Sometimes we are too young, old, ill or incapacitated—either temporarily or permanently—to be able to make an informed choice. Other times, we are at the peak of our rational powers but get muddled in discussions about risk: occasionally we form incompatible preferences based on different presentations of the same information. Tversky and Kahneman document such departures from ideal rationality.[7] These difficulties, however, do not undermine the moral obligation of clinicians to strive toward the ideal of informed choice.[8]

The moral imperative for informed choice provides good reason for considering what information is required in the ideal case. In the context of decisions regarding therapeutic

interventions the information needed for informed choice would minimally be expected to include: the therapeutic options available, their expected benefits, and possible adverse outcomes. The manner in which the possible benefits and harms of therapy are communicated is important. There is always a considerable degree of uncertainty in any science, and the clinical sciences are no different. My focus is how this uncertainty is best conveyed for adequately informing choice.

The basic intuition I wish to motivate is that the best way to communicate uncertainty regarding the possible benefits and harms of therapy is through the use of probabilities. In particular, I suggest what is required is the probability of gaining benefit or harm from a proposed therapy given the current evidence. Ideally, the communicated probabilities should be as individualised to personal circumstances as possible, relate to frequencies of events observed in medical trials (benefit or harm), and be appropriately adjusted by the clinician in light of experience in treating similar patients. This may sound both obvious, and a little idealistic. I accept both. Though, to be clear, I am not suggesting that a probabilistic discussion of benefits and harms is a necessary condition for informed choice. A rational individual may not require (nor want) this amount of detail in order to inform their decision. However, given that some individuals do[9], and the request is reasonable, the question becomes how best to move toward this ideal.

Much of the health risk communication literature talks of providing probabilities.[10] However, the "probabilities" they provide are not the probability of benefit or harm given the available evidence, experience of the clinician, and personal characteristics of the

patient. Rather, the "probabilities" conveyed in health risk communication are estimates of classically analysed trials (I provide examples in the next section). I wish to distinguish the estimates of classical statistics from the probabilities ideally suited to informing choice. The ideal input for informed choice, I argue is, the probability of benefit or harm of therapies given all available evidence. Classical statistics is unable provide this ideal because it explicitly rejects assigning probabilities to hypotheses. More on this is a moment.

In order to fix ideas, it may assist if a methodology that does assign probabilities to the benefits and harms of therapies is introduced. One such method is the Bayesian approach to statistical inference. The merits, or otherwise, of competing approaches to statistical inference is highly contentious. Rather than enter debate regarding the ideal method of statistics for science, the present focus is on the inputs classical and Bayesian approaches provide for informed choice. It is worth noting that both approaches to statistical inference play an important, and largely accepted, role sections of contemporary medical research. Assuming a degree of pluralism, it seems legitimate to pose the question of what type of probabilistic information best facilitates informed choice, independent of a favoured account of statistical inference in science.

Bayesians approach uncertainty by assigning probabilities to propositions about the world. The starting point is a prior probability. This prior probability represents current uncertainty about a proposition: say, whether drug *A* benefits a particular population. The best available evidence informs the prior. An experiment, testing drug *A* in a defined

population, provides additional evidence. This evidence is used to update the prior probability via Bayes's Theorem. The probability of the proposition, updated by the experimental evidence, is called the posterior probability. This posterior probability of the benefit of drug *A* in the defined population can be used as an input for informed choice.[11]

Bayesian's have different views on whether the probabilities used by the approach are objective or subjective. However interpreted, the Bayesian methodology provides the kind of probability I have suggested is needed to adequately inform choice. That is, a probability relativised to the patient population under consideration, related to previously observed frequencies, and reflective of the underlying theory. It may be "subjective", or better, judgemental, in the sense that clinicians may legitimately disagree on the probability of an event for a given population or individual. But is "objective", at least in the context of health care, in the sense that it should be possible to give account of the types of evidence and experience appealed to, and weighed, in providing the probability of benefit or harm. It needs to be justifiable.

In addition to the probability of an event it is also important to consider the importance or "value" associated with that event for the individual. Hence, it is useful to consider risk communication within a decision theoretic framework. Such a framework is, at least, implicit in most contemporary discussions.[12] Within this framework two components are important, the probability of the outcome (whether benefit or harm) and the utility of the outcome. For example, an individual may place similar weight in their deliberations on a

very rare, but severe adverse effect, as they do for a relatively common mild adverse effect. Probabilities and utilities inform decisions.

How probability statements that inform choice are derived and presented is my focus here. Before expanding on how classical statistics raises a problem for informed choice, I outline the contemporary approach to the communication of risk in health.

RISK COMMUNICATION IN HEALTH

A growing literature on risk communication in health takes the ethical obligation to inform choice as its foundation, and asks how this might best be achieved. Some look at "decision-analytic" tools, focussing on how utilities are most appropriately elicited, and others focus on how the probabilities might be conveyed to best inform choice.[13] Useful reviews of studies on communicating risk in health are available.[14] The effects of communicating risk in different numerical (absolute and relative risk), graphical (bar graphs displaying relative benefits and risk) and pictorial displays have been evaluated. Importantly, all of these methods use the same input for risk communication: observed estimates from randomised controlled trials, or meta-analyses.

Medical trials analysed according to classical statistics do not deal in probabilities, at least not directly. Two related aspects of classical statistics are important to our context: significance testing and estimation. With regard to significance testing, propositions about the world are considered either true or false. A medicine benefits to a specified

degree, or doesn't. It possesses a particular side effect, or not. An interpretation of probability enters classical analysis, but at no stage are probabilities directly attached to propositions about the benefits or harms of interventions. Rather, classical statistics attaches probabilities to the methods by which the results have come about. This particular interpretation of frequentist statistics considers the number of times the observed data would be expected on hypothetical repetitions of the experiment. This restricted and conceptually challenging approach to probability impedes discussion of uncertainty in clinical encounters.

Estimates from medical trials provide the quantification for risk communication. Typically, only estimates that have first passed a significance test are considered "accepted" and thus appropriate for communication to individuals.[15] To pass a significance test is to be shown to be statistically significantly different from the null hypothesis. A considerable conceptual framework underpins both hypothesis testing and estimation. There is considerable debate about this framework, even from *within* classical statistics. In particular, some argue that estimation should replace hypothesis testing.[16] Both of these approaches are used in contemporary trial analysis. Given that most readers will be familiar with the outputs of these classical approaches—that is, *p* values and 95% confidence intervals—this section will briefly illustrate how they are used in risk communication. What these outputs of classical statistics actually mean, and the problem they pose for informing choice, will be made precise in the following section.

Put simply, the risk communication literature "probabilises" the estimates of classical statistics. For example, Man-Son-Hing, et al., describe the use of a decision aid to inform individuals with atrial fibrillation who need to decide whether to take aspirin or warfarin to prevent stroke.[17] The probabilities used to inform the choice are estimates from previous trials involving aspirin and warfarin respectively.[18]

Warfarin prevents more strokes but increases the risk of bleeding. Aspirin is easier to take, and relatively safe, but is less effective than warfarin. To inform patient decisions the following estimates are provided in a range of formats: if you take aspirin for two years, you have 8% risk of a stroke and 1% risk of a severe bleed; if you take warfarin for two years, you have 4% risk of a stroke and 3% risk of a severe bleed. That these "probabilities" are all classical estimates that have "passed" a hypothesis test is what I wish to draw attention to.

Confidence intervals can be communicated as an alternative to point estimates from "accepted" hypotheses. While I am unaware of any decision aids that explicitly communicate confidence intervals rather than point estimates, there are examples where communicating confidence intervals appears more appropriate, and is, at least implicitly, urged.

The emerging controversy over the cardiovascular effects of rosiglitazone provides an example. The RECORD trial was halted early due to concern about cardiovascular events in patients taking rosiglitazone. The observed hazard ratio for hospitalisations due

to cardiac conditions or cardiac death for patients taking rosiglitazone was reported as 1.11 (95% CI, 0.93–1.32).[19] Note the absence of a *p* value. An editorial commenting on this result suggested: "the data are consistent with as much as a 7% reduction in cardiovascular risk, and as much as a 32% increase".[20] Should quantitative information on risk be provided to patients, presumably, it is these figures that should be communicated.

While these examples are not exhaustive, they are representative. Current approaches to risk communication in health do not overcome the classical statistical rejection of assigning probabilities to hypotheses; instead they "probabilise" the estimates of classical statistics.

A SKETCH OF THE CLASSICAL APPROACH TO HYPOTHESIS TESTING

The appropriate role of probability in statistical inference is controversial. Two camps in philosophy of statistics—the Bayesian, and the classical statisticians—hold widely divergent views. Whereas probabilities of hypotheses are central to Bayesian statistics, classical statisticians deny such probabilities play any useful role in scientific inference. Given that I suggest the ideal input for informed decisions is probability statements about the risks and benefits of therapies, it may not surprise that medical trials analysed according to classical statistics do not fare well in informing choice.

Understanding the classical approach to significance testing is important to informed choice, because, as discussed earlier, a benefit or harm of a drug is generally not considered "accepted" unless it has passed a significance test. A hypothesis concerning a benefit or harm of a medicine is accepted when the null hypothesis—typically the hypothesis that there is no effect of the medicine—is rejected. While the confidence interval approach adds some flexibility, in practice, most confidence intervals are only considered when a hypothesis has passed a significance test. The classical approach to testing a null hypothesis can be briefly summarised. To aid exposition, I will assume a one-sided test of a null hypothesis.

First, a way to summarise the raw data of the experiment is chosen. This is called a test statistic. It might be mortality rate over time, or something more complicated. The test-statistic provides an idealisation of the data.

Second, a sampling distribution is considered. Consider the null hypothesis that the true value of the test statistic is zero. The sampling distribution includes all the values the test statistic could take, were the null hypothesis true. It is created based on how often each possible value of the test statistic would be expected if the trial were to be repeated infinitely. Assuming, as we are, that the null hypothesis is true, we would expect that many repetitions of the experiment would result in many observed values of the test statistic being close to zero (with a reducing frequency of results for test statistics further away from zero).[21] This sampling distribution provides the primary conceptual apparatus for significance testing and the classical statistical interpretation of probability.

Third, the sampling distribution is divided into two regions; accept and reject. A pre-set "alpha level" defines the reject region. Alpha is usually (and arbitrarily) set at 0.05. It corresponds to the proportion of the sampling distribution in the tail. For a one-sided test of a null hypothesis against an alternative hypothesis of a positive effect, the alpha region will be located in the right hand tail region of the distribution. Values of the test statistic in this tail region would be expected in 5 out of every 100 repetitions of the trial.

Now the actual trial is conducted. Should a value for the test statistic that falls into the alpha region be observed, classical statistics suggests the null hypothesis can be rejected. This is because such a value for the test statistic would be expected infrequently were the null hypothesis true. Typically a *p* value is calculated for the observed test statistic. A *p* value is the proportion of the sampling distribution of the null hypothesis that corresponds to the value of the test statistic observed, or more extreme values. Hence a *p* value of 0.04 suggest that the observed test statistic, or a test statistic more extreme, would be expected only 4 out of every 100 repetitions of the experiment.

THE PROBLEM

There is more to hypothesis testing than what is described here, but what should be clear is the notion of probability at play. And in particular, how restricted this notion of probability is. Rather than contemplating the probability of an effect of a drug in a particular population, taking into account the pathophysiology of the patients, the

pharmacology of the drug, and previous evidence, classical statistics restricts itself to hypothetical repetitions of the experiment. When informing inferences, once the experiment has been designed, attention is confined to how frequently the observed value of the test statistic would be expected assuming the truth of the null hypothesis. Decision makers, individuals and clinicians, typically want to know the inverse—the probability of the hypothesis, alternative or null, based on what was observed in the trial and other relevant evidence. Classical statistics does not provide this information.

When a hypothesis concerning a risk or benefit of a medicine is accepted, attention turns to the magnitude of the effect. This is the problem of estimation. In classical point estimation, if the null hypothesis is rejected, the value of the test statistic observed in the trial is taken to be the true magnitude of the effect. If the null hypothesis is not rejected, then the true value is taken to be the value proposed by the null hypothesis. As seen in the previous section, it is this point estimate, often expressed as a probability, which is communicated to inform choice.

The confidence-interval approach relies on the same conceptual framework as hypothesis testing and point estimation. Rather than consider trial observations in light of the expectations of infinite repetitions of the trial assuming the null hypothesis, the confidence interval approach provides an interval, and considers how often the true value of the test statistic would be expected to fall in this interval, were the trial to be repeated infinitely.

For illustration, recall the confidence interval provided for the primary endpoint in RECORD: 95% CI, 0.93–1.32. What this says is that were we to repeat the trial infinitely many times we would expect the true value of the test statistic to fall within the stated interval 95 out of every 100 times. The confidence interval denotes the precision of the study. Precision, in turn, is related to the number of participants in the study: in general, the more trial participants, the higher the precision, and the narrower the confidence interval. There are benefits to confidence intervals over *p* values and hypothesis tests, but at no stage are we attaching probabilities to hypotheses about the risks and benefits of interventions.

The probabilities at play in classical statistics are difficult to grasp. They are far from ideal tools for informing choice. The confusion they cause practitioners is testament to this.[22] I do not suggest Bayesian analyses are a panacea for informed choice. Uncertainty about risks and benefits of medicines will remain, and confidence in the probabilities provided will always be constrained by the reliability of the evidence. But at least the probabilities required are available and legitimate, and they are in a communicable form.

Given their conceptual differences, it won't surprise that Bayesian statisticians can make different inferences to their classical counterparts in light of the same evidence.[23] The difference in their conception of probability is reflected in a difference in focus. Classical statistical tools are particularly focussed on the data generated in the experiment: Does the observed data differ substantially from they hypothesised null?

Prior beliefs and theoretical concerns are not explicitly considered in statistical measures such as *p* values. In contrast, Bayesian analyses can be thought to work at one level higher. Prior beliefs and theoretical concerns are explicitly considered, and incorporated into the probabilistic framework. This might be considered an advantage or a disadvantage, depending on your view on such matters. And which approach to statistical inference is preferred may depend on the scientific context. I am not entering this debate, but I am suggesting Bayesian probabilities have advantages when it comes to informing choice.

It is clear that when people talk about probability they can mean a range of things. This seems equally clear in the context of people considering their therapeutic options. Regardless of the format of inputs from clinical trials, it will always be necessary for the clinician to clarify what the probabilities being referred to are, what they are relative to, and how they have been obtained. A reply available to strict classical statisticians is to educate people. Help people become more science savvy and statistically literate. A worthy goal, particularly as most of science is analysed according to classical methods. However, given the confusion apparent even in those schooled in its methods, all options should be considered. Suggesting that the classical conception of probability is a square peg being forced into the round hole of the way we think about risk is perhaps a metaphor too far. But highlighting the awkwardness of the fit does seem appropriate.

CONCLUSION

If we are to take the ethics of informed choice seriously then we need to attend to the notions of probability at play, both in analysis of trials and communication to individuals. This claim is fairly uncontroversial, and most participants in the debate agree. Despite this, explicit discussion about differing notions of probability is rare both in the ethics and health risk communication literature (at least in the context under discussion here).

To the extent that individuals require probabilities of risks and harms in order to make a choice between therapies, the limitations of classical analyses should be recognised. Where a pluralistic approach to statistical inference is appropriate, consideration of the needs of decision makers supports the use of Bayesian methodologies.

Acknowledgements

I would like to thank Jason Grossman and Mark Colyvan for instructive discussion, and helpful comments on earlier drafts.

Notes

[1] Two immediate questions: what constraints on an individual's therapeutic choice are appropriate? And, when is 'all else equal'? It is clear, for example, that expert advice, best evidence and the availability of resources play a legitimate role in constraining therapeutic choice. It is also clear that everything is not equal when an individual's rationality is in question. Then the question then becomes: what happens to the

clinician's obligation to inform choice? For a comparison of informed choice and shared decision making in the context of general practice, see G. Elwyn, et al., "Shared Decision Making And The Concept Of Equipoise: The Competencies Of Involving Patients In Healthcare Choices," *Br J Gen Pract* 50 (2000): 892–899. For one example of a discussion on informed consent when rationality is in question, see: J. Savulescu & R. Momeyer, "Should Informed Consent Be Based On Rational Beliefs?" *J Med Ethics* 23 (1997): 282–288.

[2] A. Tversky & D. Kahneman, "Judgement Under Uncertainty: Heuristics and Biases," *Science* 185 (1974): 1124–1131; R.M. Epstein, et al., "Communicating evidence for participatory decision making," *JAMA* 291 (2004): 2359–2366.

[3] World Medical Association, 1964 (Last updated 2004). *World Medical Association Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects.*

[4] R.M. Veatch, "Abandoning Informed Consent," *Hastings Cent Rep* 25 (1995): 5–12.

[5] L. Doyal, "Informed Consent: Moral Necessity Or Illusion?" *Qual Health Care* 10 (2001): i29–i33.

[6] O. O'Neill, "Some Limits of Informed Consent," *J Med Ethics* 29 (2003): 4-7.

[7] For a general discussion, see A. Tversky & D. Kahneman, cited in n. 2 above; for a discussion in relation to health care, see Epstein, et al., also cited in n. 2 above.

[8] L. Doyal, cited in n. 5 above; M. Kottow, "The battering of informed consent," *J Med Ethics* 30 (2004): 565–569.

[9] S. Ford, T. Schofield and T. Hope, "What are the ingredients for a successful evidence based patient choice consultation?: A qualitative study," *Social Science and Medicine* 56 (2003): 589–602.

[10] Epstein, et al. cited in note n. 2 above; M. Man-Son-Hing, et al., "A Patient Decision Aid Regarding Antithrombotic Therapy for Stroke Prevention in Atrial Fibrillation: A Randomised Controlled Trial," *JAMA* 282 (1999): 737–743.

[11] For a review of Bayesian analysis in clinical trials, see D.A. Berry, "Bayesian Clinical Trials," *Nat Rev Drug Discov* 5 (2006): 27.

[12] A. Edwards, and G. Elwyn, "Understanding Risk And Lessons For Clinical Communication About Treatment Preferences," *Qual Health Care* 10 (2001): i9-i13.

[13] N.F. Col, et al., "Short-term Menopausal Hormone Therapy for Symptom Relief: An Updated Decision Model," *Arch Intern Med* 164 (2004): 1634-1640; provides an example of eliciting "utilities". For an example of risk communication, see Man-Son-Hing, et al., cited in n. 10 above.

[14] See, for example Epstein, et al., cited in n. 2 above; A. Edwards, et al., "Explaining Risks: Turning Numerical Data Into Meaningful pictures," *BMJ* 324 (2002): 827-830.

[15] I put "accept" in scare quotes to acknowledge that classical statistics never fully accepts a hypothesis that has passed a test. This is in line with Popperian philosophy of science. Classical statistics provides methods for deciding when a hypothesis, usually a "null hypothesis", can be rejected. It is when the conditions for rejection of the null hypothesis are appropriately met, that the alternative hypothesis can be provisionally "accepted".

[16] J.H. Ware, et al., "P Values." in *Medical Uses of Statistics.* J.C.I. Bailar & F. Mosteller, eds., (NEJM Books, Boston 2002): 480.

[17] Man-Son-Hing, et al., cited in n. 10 above.

[18] The Atrial Fibrillation Investigators, "Risk factors for stroke and efficacy of antithrombotic therapy in atrial fibrillation," *Arch Intern Med* 154 (1994): 1449–1457; The SPAF III Writing Committee for the Stroke Prevention in Atrial Fibrillation Investigators, "Patients with nonvalvular atrial fibrillation at low risk of stroke during treatment with aspirin," *JAMA* 279 (1998): 1273–1277.

[19] P.D. Home, et al., "Rosiglitazone Evaluated for Cardiovascular Outcomes—An Interim Analysis," *N Engl J Med* 357 (2007): 28–38.

[20] J.M. Drazen, et al., "Rosiglitazone—Continued Uncertainty about Safet," *N Engl J Med* 357 (2007): 63-64.

[21] This also assumes that the sampling distribution is monotonic, and assumes that any measurement error is small and distributed randomly.

[22] C. Poole, "Low P-Values or Narrow Confidence Intervals: Which are more durable?" *Epidemiology* 12 (2001): 291-294; J.M. Young, et al., "General Practitioners' Self Ratings Of Skills In Evidence Based Medicine: Validation Study," *BMJ* 324 (2002): 950-951.

[23] J. Berger & T. Sellke, "Testing A Point Null Hypothesis: The Irreconcilability Of P-Values And Evidence," *J Am Stat Assoc* 82 (1987): 112-122; Berry, cited in n. 11 above.

Bibliograhy

Association, World Medical. "World Medical Association Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects." 1964 (Last updated 2004).

Berger, JO, and T Sellke. "Testing a Point Null Hypothesis: The Irreconcilability of P-Values and Evidence." *Journal of the American Statistical Association* 82, no. 397 (1987): 112-22.

Berry, Donald A. "Bayesian Clinical Trials." *Nature Reviews Drug Discovery* 5, no. 1 (2006): 27.

Col, Nananda F, Griffin Weber, Anne Stiggelbout, John Chuo, Ralph D'Agostino, and Phaedra Corso. "Short-Term Menopausal Hormone Therapy for Symptom Relief: An Updated Decision Model." *Archives of Internal Medicine* 164 (2004): 1634-40.

Doyal, L. "Informed Consent: Moral Necessity or Illusion?" *Quality in Health Care* 10, no. Supplement I (2001): i29-i33.

Drazen, Jeffrey M., Stephen Morrissey, and Gregory D. Curfman. "Rosiglitazone---Continued Uncertainty About Safety." *New England Journal of Medicine* 357, no. 1 (2007): 63-64.

Edwards, A, and G Elwyn. "Understanding Risk and Lessons for Clinical Communication About Treatment Preferences." *Quality in Health Care* 10, no. Supplement I (2001): i9-i13.

Edwards, A., G. Elwyn, and A. Mulley. "Explaining Risks: Turning Numerical Data into Meaningful Pictures." *Bmj* 324, no. 7341 (2002): 827-30.

Elwyn, G, A Edwards, P Kinnersley, and R Grol. "Shared Decision Making and the Concept of Equipoise: The Competencies of Involving Patients in Healthcare Choices." *British Journal of General Practice* 50, no. 460 (2000): 892-9.

Epstein, Ronald M., Brian S. Alper, and Timothy E. Quill. "Communicating Evidence for Participatory Decision Making." *Journal of American Medical Association* 291, no. 19 (2004): 2359-66.

Ford, Sarah, Theo Schofield, and Tony Hope. "What Are the Ingredients for a Successful Evidence Based Patient Choice Consultation? A Qualitative Study." *Social Science & Medicine* 56 (2003): 589-602.

Home, Philip D., Stuart J. Pocock, Henning Beck-Nielsen, Ramon Gomis, Markolf Hanefeld, Nigel P. Jones, Michel Komajda, John J. V. McMurray, and Record Study Group the. "Rosiglitazone Evaluated for Cardiovascular Outcomes -- an Interim Analysis." *New England Journal of Medicine* 357, no. 1 (2007): 28-38.

Investigators, Atrial Fibrillation. "Risk Factors for Stroke and Efficacy of Antithrombotic Therapy in Atrial Fibrillation." *Archives of Internal Medicine* 154 (1994): 1449-57.

Investigators, The SPAF III Writing Committee for the Stroke Prevention in Atrial Fibrillation. "Patients with Nonvalvular Atrial Fibrillation at Low Risk of Stroke During Treatment with Aspirin." *Journal of American Medical Association* 279 (1998): 1273-7.

Kottow, M. "The Battering of Informed Consent." *J Med Ethics* 30, no. 6 (2004): 565-69.

Man-Son-Hing, Malcolm, Andreas Laupacis, Annette O'Connor, Jennifer Biggs, Elizabeth Drake, Elizabeth Yetisir, and Robert Hart. "A Patient Decision Aid Regarding Antithrombotic Therapy for Stroke Prevention in Atrial Fibrillation: A Randomised Controlled Trial." *Journal of American Medical Association* 282, no. 8 (1999): 737-43.

O'Neill, O. "Some Limits of Informed Consent." *J Med Ethics* 29, no. 1 (2003): 4-7.

Poole, Charles. "Low P-Values or Narrow Confidence Intervals: Which Are More Durable?" *Epidemiology* 12, no. 3 (2001): 291-4.

Savulescu, J, and RW Momeyer. "Should Informed Consent Be Based on Rational Beliefs?" *Journal of Medical Ethics* 23, no. 5 (1997): 282-8.

Tversky, Amos, and Daniel Kahneman. "Judgement under Uncertainty: Heuristics and Biases." *Science* 185, no. 4157 (1974): 1124-31.

Veatch, Robert M. "Abandoning Informed Consent." *The Hastings Center Report* 25, no. 2 (1995): 5-12.

Ware, James H., Frederick Mosteller, Fernando Delgado, Christl Donnelly, and Joseph A Ingelfinger. "P Values." In *Medical Uses of Statistics*, edited by John C. III Bailar and Frederick Mosteller, 480. Boston: NEJM Books, 1992.

Young, J. M., P. Glasziou, and J. E. Ward. "General Practitioners' Self Ratings of Skills in Evidence Based Medicine: Validation Study." *British Medical Journal* 324, no. 7343 (2002): 950-1.