European Journal of Epidemiology (2005) 20: 739–745
DOI 10.1007/s10654-005-2028-1

METHODS

Molecular bias

John P.A. Ioannidis
Department of Hygiene and Epidemiology, University of Ioannina School of Medicine, Ioannina , Greece

Accepted in revised form 15 August 2005

Abstract. Bias is ubiquitous in research. The advent of the molecular era provides a unique opportunity to study the consequences of bias with large-scale empirical evidence accumulated in the massive data produced by the current discovery-oriented scientific effort, rather than just with theoretical speculations and constructs. Here I discuss some empirical evidence about manifestations of bias in molecular epidemiology. Bias may manifest as either heterogeneity or as deviation from the true estimates. The failure to translate molecular knowledge and the failure to replicate information are some typical hallmarks of bias at action. The acquired knowledge about the behaviour and manifestations of bias in molecular fields can be transferred back also to more traditional fields of epidemiology and medical research. Getting rid of false claims of the past is at least as important as producing new scientific discoveries. In many fields, the observed effects sizes that circulate as established knowledge are practically estimating only the net bias that has operated in the field all along. Issues of plausibility (in particular biological plausibility), replication, and credibility that form the theoretical basis of epidemiology and etiological inference can now be approached with large-scale empirical data.

Key words: Bias, Replication, Validity, Research, Molecular Medicine

Introduction

Bias may be broadly defined as any deviation from the truth. It could involve the distortion, negation, or suppression of the truth. The ''truth'' is by default unknown and at best being sought, so what we see and measure is the composite of the truth and various biases. Bias in scientific measurement can be conscious, subconscious, or unconscious. It can affect single measurements, sets of measurements, datasets, whole studies, sets of studies, scientific domains, and science across many domains. By extrapolation, it can also penetrate human activities beyond scientific limits. One may create theories about bias or may study its consequences. Epidemiology has been the scientific discipline par excellence to be concerned with theories of bias. We have solid and robust theories on confounding, misclassification, and information biases. We have long lists of different sorts of biases [1]. However, beyond theory, it is the consequences of bias that we measure, witness, and eventually suffer. In biomedicine in particular, not only scientists, but also people (both healthy people and patients) truly suffer unfortunately from the consequences of bias. Besides theory, it would be greatly informative to learn from bias empirically.

''Molecular bias'' may be a provocative title. I use this term here to discuss the manifestations of bias in the newer molecular era of biomedicine. What is special, if anything, about bias in the molecular era? First, there is an exponential increase in the available data with apocalyptic promises of bio-information. Second, we increasingly recognize that we are dealing with quanta of mostly small-effect risk factors. Third, we are increasingly adopting discovery-oriented, hypothesis-free approaches trying to accommodate the capabilities of massive data production. Under these circumstances where tons of data is available, empirical data on bias become very wealthy and can more than match theoretical considerations in giving us insights about how bias operates and manifests itself. Thus, the new age offers a perfect opportunity of measuring and witnessing, and hopefully not suffering, the consequences of bias.

Bias at action

Bias may arise from a huge variety of sources, even for the simplest scientific question. Bias components interact and may cancel out among themselves or strengthen each other. The composite bias may eventually manifest in two ways: as heterogeneity in

*This paper is based on a keynote lecture given at the annual graduation meeting of the Department of Epidemiology and Biostatistics of the Erasmus University of Rotterdam in June 22, 2005.

the results of variously biased studies; or as a net deviation in the results of variously biased studies and their summaries. For those initiated in evidence-based medicine, it is implicit that the typical example of a summary exercise is provided by a meta-analysis [2, 3].

Heterogeneity due to bias may be confused with genuine biological and etiological diversity. On the other hand, net diversion from the true estimates cannot and should not be confused. The problem is that we never know the true estimates, we never have a perfect gold standard for what is true. Thus, one needs to come up with some operative rules to detect bias through its consequences.

These rules cannot be perfect and they need to be continuously appraised, revised, and refined with more evidence. However, one major hint for the presence of bias may be that something does not work, as it was supposed to work based on early data, evidence and reasonable expectations.

Such "bias at action" is exemplified by the failure of translational efforts for major basic science discoveries [4, 5]. Medical progress depends on important discoveries from basic biomedical research. Among 101 highly-promising basic science publications that made clear promises in 1979–1983 that they would have imminent clinical or public health translation, only 27 had an RCT published within the next 20–25 years, 19 had a "positive" RCT, and only 5 are currently in clinical use [4]. Of these five, only 1 really has made a large clinical impact [4].

This translation failure in the communication of basic, preclinical, clinical and public health sciences shows composite bias operating at large and at a broad range. Bias should not be seen narrowly as erroneous measurements or exaggerated studies. As I mentioned above, it also incorporates the suppression of the truth, and the negation of the truth. Some of these discoveries may have been simply based on wrong or exaggerated data and inferences. However, others may have been true but never managed to translate themselves to something useful and tangible. Interestingly, the strongest predictor for the translation of basic science discoveries to clinical trials in this empirical evaluation was the involvement of the industry [4]. We already have a hint for some biased selection here on what moves forward in science.

Major postulated problems of molecular research

Research in the molecular era has to meet several challenges. There is a huge number of biological factors that can now be targeted by research efforts [6, 7]. Their effect sizes are likely to be small or very small. Sample sizes have not increased considerably compared with the past, while the complexity of the research questions has. Case–control studies of a couple of dozen or a couple of 100 people are still the norm, while a typical experiment may allow testing hundreds and thousands of "risk factors" with effect sizes that would require thousands of subjects to have adequate power for detecting each of them alone. At the same time, the old epidemiology problems of confounding and misclassification are still operating in the molecular era.

All of these problems and limitations may eventually manifest themselves as lack of replication validity in the proposed molecular findings [7, 8]. However, one of the great advantages of the molecular revolution is the ability to study bias with large-scale data. This can be useful in itself. Moreover, it can also offer insights on how bias may operate also in more traditional clinical and epidemiological disciplines.

A typical example is studies assessing genetic risk factors. There are over 10,000 publications of such studies already (e.g. one may examine the accumulated database at the Human Genome Epidemiology Network website [9] at www.cdc.gov/genomics/hugenet) and an unknown number of unpublished studies. The number as well as the complexity of targeted risk factors increases rapidly. Median sample sizes are in the range of 250 subjects though [10]. It should not be surprising thus that many of the postulated genetic effects for complex common diseases are refuted by subsequent evidence. In fact, a common scenario is to observe a large genetic effect in the first study that is nevertheless dissipated gradually towards the null as more data accumulate from additional studies on the same association [8]. Early findings of epidemiological associations in molecular genetics apparently have no predictive ability against the findings of subsequent research on the same associations [11].

Diminishing effects and refutations in traditional research disciplines

This paradigm of diminishing effect sizes can give us also insight for more traditional research disciplines. Non-replicated diminishing effects can occur also with traditional research designs, including randomized controlled trials that have typically been ranked as high-level evidence in making recommendations. An empirical evaluation of 100 meta-analyses of mental health related interventions, shows that for pharmacotherapies, it has been far more likely for effect sizes to diminish rather than increase with the appearance of newer trials [12]. Even the most influential clinical research published in the top journals may often be refuted [13]. Among highly-cited articles published in 1990–2003 and receiving over 1000 citations by 2004, 5 of 6 efficacy findings based on non-randomized trials were already contradicted or found to be exaggerated by 2004. Moreover, efficacy

findings were already contradicted or found to be exaggerated in 9 of 39 interventions [13]. Many of these contradictions such as the impact of various vitamins on cardiovascular or cancer mortality [14], or the effect of hormone replacement therapy on coronary artery disease in postmenopausal women [15] have stirred extensive debates.

Contradictions and refutations have been seen as something extravagant and unnatural that is shaking the foundations of epidemiology and clinical medicine. In fact, contradictions and refutations are simply the natural course of science. Instead of demonizing their catastrophic effects, we should embrace them as great opportunities to learn about how bias is operating and manifesting itself. In this regard, molecular epidemiology can give us 100- or 1000-fold more examples of contradictions and refutations that traditional epidemiology and clinical research has done to-date. We have lots to learn from these data.

Refutations provide empirical credibility to the notion that any epidemiologic or other association is tentative and science is a work-in-progress. Large scale testing can give us empirical evidence not only about postulated associations that are refuted. Conversely, in selected cases, associations may eventually be proven to be important despite the fact that early studies had provided inconclusive or even seemingly null results [8].

How many analyses?

Three years ago, in a naïve calculation, I estimated that there are 100 trillion possible analyses in the field of gene-disease associations alone [9]. Today, I would need to update this estimate by increasing it by at least 10 fold. This is only one circumscribed field of molecular epidemiology. The total number of epidemiological analyses that can be undertaken in much higher, possibly in the range of many quadrillions or even quintillions. Moreover, this number is increasing rapidly, as our ability to measure biological parameters improves and scientific efforts expand.

Indirect evidence and plausibility

Molecular science also provides a great opportunity to address empirically what is the predictive value of indirect evidence. By indirect evidence I mean evidence not stemming from the data and analyses-at-hand, but from other lines of inference. In any epidemiologic question that can be posed in the molecular era, there is also the possibility to amass biological and functional information about its overall significance. How pertinent is this information?

Scientists always make inferences based not only on the data at hand but also by corroborating from other lines of evidence, not only from their field, but also from different fields. Looking back at the Discussion sections of some old epidemiology papers about traditional risk factors, these other lines of reasoning may sound incongruent at present. Much of the supporting external evidence on the plausibility of the epidemiological findings has been abandoned, refuted, or dismissed. In particular "biological plausibility" becomes key in molecular epidemiology, since the probed associations in the molecular era are by definition closer to the biology of health and disease – in theory at least. So is it going to be better now that we come "closer" to the biology?

Biological plausibility can be a fertile ground for bias. We now have an excellent opportunity to study this at a large-scale. Scientists may invoke biological plausibility for almost any finding and any association, especially if this is done after the fact. Just in the year 2002 studies were published addressing the relationship of the APOE epsilon polymorphism [9] with about 40 different diseases and conditions, many of which had very little in common. It is tough to presume that truly there is some valid plausibility for all, or even a few, of these postulated associations. Biological plausibility seems to be the aspect where the imagination of scientists works at its best. In this aspect, it can be very creative. At the same time, it may be like a room of distorting mirrors: even if you place a simple brick in the middle, you can see dozens of monsters depicted all around. This can easily create and sustain a scientific literature of phantoms.

We need to amass data on the extent of this bias and on how much we can gain from indirect evidence [16]. Using our large accumulated databases, we may be able to address questions about the positive and negative likelihood ratios conferred by various lines of indirect evidence, especially biological evidence [17].

Ubiquitous and topic-specific indicators of bias

As I mentioned upfront, bias may present in many different ways and often the most important ones to consider are specific to the scientific question. Unfortunately, there is no perfect recipe about detecting bias efficiently. However, one could think of some ubiquitous indicators that might be useful to probe in a routine basis across molecular evidence, as it accumulates on a specific question.

First, bias may produce heterogeneity in the results of different studies on the same research questions. This will be tough to differentiate from heterogeneity that is due to genuine variability, but it is worthwhile documenting heterogeneity routinely as evidence accumulates. This is a lesson that has been learnt already from traditional medical research, including clinical trials [18, 19].

Second, bias may produce different effects in studies of various characteristics that make them less

or more susceptible to biases. There are two parameters that can always be considered in any type of epidemiological study regardless of the question that it addresses. These are the mass of the evidence and the time of the evidence. One may always examine whether small studies (or studies with greater imprecision in their effect estimates) differ in their results from larger studies. The presence of such differences has been often translated as synonymous of publication bias [20]. This is a gross misrepresentation [19]. Such differences could represent any type of bias, or, again, any type of genuine heterogeneity. Regarding time, one may always examine whether the early studies on a specific research question differ from the results of subsequent studies [8, 21, 22]. One may evaluate whether there are changes in the total evidence over time [8, 21, 22]. Science is cumulative and changes in the evidence are suggestive of bias – or, again, genuine heterogeneity.

An empirical evaluation of gene-disease association studies shows that although it is very common to detect significant gene-disease associations when several studies accumulate, in the large majority of these cases, there are also hints of heterogeneity or bias: study results differ beyond chance, larger studies disagree with smaller ones, and/or later studies disagree with earlier ones [23]. This means that these associations are often underlying heterogeneity or bias. Interestingly, discrepancies in time are relatively independent of discrepancies in mass, suggesting that these two parameters offer complementary views.

Successive extremes

The Proteus phenomenon [24] is another interesting behaviour of the accumulating molecular evidence. In molecular epidemiology, it is very common to see an early succession of the most extreme opposing estimates of effect followed by studies that show intermediate results. Typically, the study with the most prominent effect size is published first. This is immediately followed by the study that shows the most opposite effect ever observed. This creates extreme between-study heterogeneity in the results of these early studies. Many studies then follow that report results that are intermediate between these two extremes. The between-study variance diminishes and the summary data across all studies seem to gravitate towards some consensus.

The rapid succession of extreme opposites may be due to the fact that molecular epidemiology results can be generated very rapidly currently. Once a new epidemiological association in proposed, many investigators can try to replicate it extremely rapidly. Just as the team that finds the strongest association may have an advantage to publish its results first, the team that finds the most contradictory finding may also have an advantage for rapid publication. It could even be that the contradictory data are already available in some cases. This is a speculation of course, but given the massive information produced currently, this may become a common situation, unless a mechanism is set to allow the transparent dissemination of all produced results. According to this scenario, the contradictory results have already been obtained in the past, but their investigators don't deem them important to disseminate. Then the extreme association gets published. The contradictory data assume value due to their contradictory nature and thus they gain priority for writing up for publication and for getting published!

When the future has already happened

Under these circumstances, a dreadful evolution of science may be envisioned where many "new" scientific theories have been already tested and rejected in the past – but simply this is not widely known. This becomes known only when someone claims to have found something. The future has already happened. Is this just science fiction? Even if in the extreme form, this is probably not true (until now at least), we have increasing evidence that scientific data are not always visible: there are readily available, available, hidden, very well hidden, and disappeared data. The process by which a piece of scientific data finds itself in each of these categories can be complex and topic-specific, but the presence of these categories of visibility is becoming increasingly apparent in some fields such as the literature of molecular prognostic factors [25].

Effect size = net bias

Scientists and the public are typically interested on what is likely to be true, what association is important, what treatment works, etc. Certainly this is important to peruse, but it can also be illuminating to estimate what is not true and how much the net bias is. For research fields where there is a very intense research activity on thousands, millions, and billions of information items, but only a few of them lead to fruitful "true" discoveries, the vast majority of the findings will be false positive. Consider radio astronomy signals in the search for extraterrestrial intelligence. To-date all the signals received are considered to be noise (bias) due to various artefacts. If there is no extraterrestrial intelligence after all, all the signals would be simple bias. Even if extraterrestrial intelligence does exist, and a couple of signals eventually turn out to be true leads, still the very vast majority of the signals would simply be bias.

Back to epidemiology and biomedicine, in fields with extremely large amounts of information and no

or few true positives (''needles in a haystack''), the effect sizes of all the identified false positive findings are estimates of the net bias that has been exercised in each one of them [26]. The distribution of these effect sizes provides an accurate description of the magnitude and behaviour of the net bias in the field at large. The few true findings will be negligible within this sea of false positives. Under such circumstances, if the expected effect size is denoted by θ and there are n biases θ_i ($i = 1,\ldots, n$) operating in the same or different directions (some increasing and some decreasing the observed effect), it is fair to say [26] that

$$\theta = \Sigma\varepsilon_i$$

or in simple words

Effect size = net bias

Thus, disciplines that find larger effect sizes (those that incidentally are scientifically considered more successful) are simply more biased than others that find smaller effect sizes. Similarly, in the same scientific discipline, the most successful and appreciated studies are simply the ones that suffer the worst net bias. Furthermore, as I have shown previously, it can be proven [26] that the post-study odds of a true finding are small:

- When effect sizes are small
- When studies are small
- When the field is ''hot'' (many teams work on it)
- When there is strong interest in the results
- When databases are large
- When analyses are more flexible

To-date most biomedical studies have been small anyhow and most effect sizes that we chase in epidemiology and biomedicine are also small. Besides these two characteristics, all the other characteristics listed above are all attributes of what would be considered an active, interesting, stimulating, promising scientific field nowadays. This is largely a faithful picture of modern science.

Autonomy vs. registration, and the pyramid of credibility

Much of the problem arises from the very advantages of the best mode of scientific thinking. Science at its best is highly creative and independent, even fiercely individualistic and competitive. Scientific thinking often tends to be selective, flexible, even undisciplined, in the quest for meaningful information. Efforts for transparency, discipline, and structured, comprehensive reporting of information may seem to go against this blessed autonomy. It is not possible to put all scientific information into stereotyped boxes, this could be devastating to free thought. However, there must be some way to keep track of what is happening.

Upfront study registration has been adopted for randomized clinical trials, as a means for minimizing publication and reporting biases and maximizing transparency [27]. This is very reasonable, and hopefully it will have a beneficial effect on clinical research and its credibility. Clinical trial investigators cannot complain that their ideas are stolen by being forced to say upfront that they are doing a study and convey to fellow scientists how they are doing it and what they expect to measure and how. Some thinkers have reached the point to claim that clinical trials should not even published in scientific journals; they should only be registered fully initially as protocols, and, when finished, along with their results [28].

However, for molecular research, upfront registration in public of all ideas is counter intuitive and goes against the individualistic spirit of discovery in basic research. Nevertheless, instead one could aim for registries of investigators and data specimen collections [29]. These could be inclusive networks of investigators working on the same disease, field and set of molecular factors. These networks could promote better methods and standardization in their field, maintain the research freedom for individual participating teams and help create a mentality where there would be thorough and unbiased testing of proposed hypotheses with promising preliminary data on large-scale comprehensive databases [29]. Due credit would be given to investigators for both ''positive'' and ''negative'' findings and all data would be transparently available. There are already several examples of such comprehensive networks and we need to peruse further this mode of operation [29, 30].

As data accumulate on molecular research questions, we should also increasingly acknowledge that they are only providing a work-in-progress. At any time point, a large portion, perhaps the majority, or even the vast majority, of proposed or even (tentatively) accepted findings in a research field will simply be false positives, the result of bias [26, 31]. We need to learn not only from the big successes, but primarily from the big failures of discovery-oriented research, for example in microarrays or proteomics [32–35].

A few of these findings will be replicated though and this should lead to further and further replication until they reach a high enough credibility level. For questions where humans are put at risk to answer them, equipoise for expected benefits vs. risks to humans will need to be considered each time an effort is designed to improve the credibility level. However, for most research questions in the molecular era, this point of concern where this kind of equipoise is violated is unlikely to be reached for the majority of research questions.

It may end up being mostly an issue of resources and allocation of scientific priority. As the potential for measuring increases geometrically, there will be competition on whether one should focus on getting more and more new measurements that have become feasible or trying to replicate again and establish

some of the older observations. To-date, scientific circles, biomedical journals, grants agencies, and even the lay public strongly support flash innovation rather than careful confirmation. If this spirit continues to prevail, we will end up with tons of non-replicated and often even unchallenged findings forming the corpus of a very unreliable ''science''.

We need to learn to live with tons of bias floating around, simply we should not put too much credibility on these pieces of information, but take them for what they are: interesting, preliminary, evolving, in-progress, challenging, stimulating investigations, most likely closer to bias than to the truth. The credibility pyramid already has a large basis of findings of low credibility and a narrow top of findings with high credibility [31]. The base of the pyramid will probably continue to become broader and broader.

Closing points

I hope I have made a convincing (even if biased!) point that whatever we measure may be biased. Bias is ubiquitous. This is fine, provided we recognize the problem and grasp this fantastic, energizing opportunity to work with bias and to study it. Scientific findings should be ascribed a credibility level that is different from and goes beyond their formal statistical significance from statistical hypothesis testing [26]. In the past, we had few research findings, while currently we have too many research findings. Therefore, getting rid of tentative-but-wrong research findings should become at least as important as finding new ones. Finally, the new era can also offer us insights on the credibility of findings in traditional epidemiology disciplines where bias was mostly the focus of theoretical speculation until now.

References

1. Sackett DL. Bias in analytic research. J Chronic Dis 1979; 32: 51–63.
2. Olkin I. Meta-analysis: Reconciling the results of independent studies. Stat Med 1995; 14: 457–472.
3. Lau J, Ioannidis JP, Schmid CH. Quantitative synthesis in systematic reviews. Ann Intern Med 1997; 127: 820–826.
4. Contopoulos-Ioannidis DG, Ntzani E, Ioannidis JP. Translation of highly promising basic science research into clinical applications. Am J Med 2003; 114: 477–484.
5. Ioannidis JP. Materializing research promises: Opportunities, priorities and conflicts in translational medicine. J Transl Med 2004; 2(1): 5.
6. Ransohoff DF. Bias as a threat to the validity of cancer molecular-marker research. Nat Rev Cancer 2005; 5: 142–149.
7. Ransohoff DF. Rules of evidence for cancer molecular-marker discovery and validation. Nat Rev Cancer 2004; 4: 309–314.
8. Ioannidis JP, Ntzani EE, Trikalinos TA, Contopoulos-Ioannidis DG. Replication validity of genetic association studies. Nat Genet 2001; 29: 306–309.
9. Khoury MJ, Little J. Human genome epidemiologic reviews: The beginning of something HuGE. Am J Epidemiol 2000; 151: 2–3.
10. Ioannidis JP. Genetic associations: False or true? Trends Mol Med 2003; 9: 135–138.
11. Trikalinos TA, Ntzani EE, Contopoulos-Ioannidis DG, Ioannidis JP. Establishment of genetic associations for complex diseases is independent of early study findings. Eur J Hum Genet 2004; 12: 762–769.
12. Trikalinos TA, Churchill R, Ferri M, et al. Effect sizes in cumulative meta-analyses of mental health randomized trials evolved over time. J Clin Epidemiol. 2004; 57: 1124–1130.
13. Ioannidis JP. Contradicted and initially stronger effects in highly cited clinical research. JAMA 2005; 294: 218–228.
14. Lawlor DA, Davey Smith G, Kundu D, Bruckdorfer KR, Ebrahim S. Those confounded vitamins: What can we learn from the differences between observational versus randomised trial evidence? Lancet 2004; 363: 1724–1727.
15. Rossouw JE, Anderson GL, Prentice RL, et al. Risks and benefits of estrogen plus progestin in healthy postmenopausal women: Principal results From the Women's Health Initiative randomized controlled trial. JAMA 2002; 288: 321–333.
16. Wacholder S, Chanock S, Garcia-Closas M, El Ghormli L, Rothman N. Assessing the probability that a positive report is false: An approach for molecular epidemiology studies. J Natl Cancer Inst 2004; 96: 434–442.
17. Rebbeck TR, Spitz M, Wu X. Assessing the function of genetic variants in candidate gene association studies. Nat Rev Genet 2004; 5: 589–597.
18. Lau J, Ioannidis JP, Schmid CH. Summing up evidence: One answer is not always enough. Lancet 1998; 351: 123–127.
19. Ioannidis JPA. Differentiating biases from genuine heterogeneity: Distinguishing artifactual from substantive effects. In: Rothstein HR, Sutton AJ, Borenstein M (eds), Publication bias in meta-analysis: Prevention, assessment and adjustments. John Wiley and Sons, New York: 2005.
20. Egger M, Davey Smith G, Schneider M, Minder C. Bias in meta-analysis detected by a simple, graphical test. BMJ 1997; 315: 629–634.
21. Ioannidis J, Lau J. Evolution of treatment effects over time: Empirical insight from recursive cumulative metaanalyses. Proc Natl Acad Sci USA 2001; 98: 831–836.
22. Ioannidis JP, Contopoulos-Ioannidis DG, Lau J. Recursive cumulative meta-analysis: A diagnostic for the evolution of total randomized evidence from group and individual patient data. J Clin Epidemiol 1999; 52: 281–291.
23. Ioannidis JP, Trikalinos TA, Ntzani EE, Contopoulos-Ioannidis DG. Genetic associations in large versus small studies: An empirical assessment. Lancet 2003; 361: 567–571.
24. Ioannidis JP, Trikalinos TA. Early extreme contradictory estimates may appear in published research: The Proteus phenomenon in molecular genetics

research and randomized trials. J Clin Epidemiol 2005; 58: 543–549.
25. Kyzas PA, Loizou KT, Ioannidis JP. Selective reporting biases in cancer prognostic factor studies. J Natl Cancer Inst 2005; 97: 1043–1055.
26. Ioannidis JP. Why most published research findings are false. PLoS Med 2005; 2(8): e124.
27. De Angelis CD, Drazen JM, Frizelle FA, et al. Is this clinical trial fully registered? A statement from the International Committee of Medical Journal Editors. Ann Intern Med 2005; 143: 146–148.
28. Smith R. Medical journals are an extension of the marketing arm of pharmaceutical companies. PLoS Med 2005; 2(5): e138.
29. Ioannidis JP, Bernstein J, Boffetta P, et al. A network of investigator networks in human genome epidemiology. Am J Epidemiol 2005; 162: 302–304.
30. Ioannidis JP, Rosenberg PS, Goedert JJ, O'Brien TR. International meta-analysis of HIV host genetics. Commentary: Meta-analysis of individual participants' data in genetic epidemiology. Am J Epidemiol 2002; 156: 204–210.
31. Ioannidis JPA. Grading the credibility of molecular evidence for complex diseases. Int J Epidemiol 2005 (in press).
32. Michiels S, Koscielny S, Hill C. Prediction of cancer outcome with microarrays: A multiple random validation strategy. Lancet 2005; 365: 488–492.
33. Ioannidis JP. Microarrays and molecular research: Noise discovery? Lancet. 2005; 365: 454–455.
34. Ntzani EE, Ioannidis JP. Predictive ability of DNA microarrays for cancer outcomes and correlates: An empirical assessment. Lancet 2003; 362: 1439–1444.
35. Ransohoff DF. Lessons from controversy: Ovarian cancer screening and serum proteomics. J Natl Cancer Inst 2005; 97: 315–319.

Address for correspondence: John P.A. Ioannidis, MD, Department of Hygiene and Epidemiology, University of Ioannina School of Medicine, Ioannina 45110, Greece
Phone: +302651097807; Fax: +302651097867
E-mail: jioannid@cc.uoi.gr