

Viewpoint: Replication in economics

Daniel S. Hamermesh *Department of Economics,
University of Texas at Austin*

Abstract. This examination of the role and potential for replication in economics points out the paucity of both pure replication – checking on others’ published papers using their data – and scientific replication – using data representing different populations in one’s own work or in a comment. Several controversies in empirical economics are used to illustrate how and how not to behave when replicating others’ work. The incentives for replication are examined, and proposals aimed at journal editors and authors are advanced that might stimulate an activity that most economists applaud but few perform. JEL classification: B41, A14, C59

Réplication en science économique. L’examen du rôle potentiel de la réplication en science économique souligne la pauvreté à la fois de la réplication pure (la vérification des résultats des textes publiés en utilisant les données utilisées dans ces travaux) et de la réplication scientifique (l’utilisation de données tirées de populations différentes pour valider des résultats). Plusieurs controverses en économie empirique sont utilisées pour illustrer comment se comporter et ne pas se comporter dans la réplication des travaux d’autrui. Les incitations à faire de la réplication sont analysées, et certaines suggestions sont faites aux rédacteurs de revues et aux auteurs afin de stimuler une activité plus grande dans ce registre – que les économistes applaudissent mais que bien peu acceptent de pratiquer.

1. Introduction: What is replication?

Economists treat replication the way teenagers treat chastity – as an ideal to be professed but not to be practised. Why is this? How much replication is done by

Edward Everett Hale Centennial Professor of Economics, University of Texas at Austin, and research associate, NBER and IZA. I thank Bernd Fitzenberger, Bruce McCullough, and Gerald Oettinger for helpful suggestions on an earlier draft, the authors of several of the studies cited here for useful clarifications of their views on the controversies in which they were involved, and several editors for their experiences at their journals. Email: hamermes@eco.utexas.edu

economic researchers? What are the incentives/disincentives to engage in replication exercises? These are positive questions. Does our current treatment of replication lead to credible knowledge about human behaviour? Does it do so in a manner that is socially optimal? These are normative questions.

Even before these questions can be answered, we need to clarify what we mean by replication in the context of economic research. One dictionary defines replicate as ‘duplicate, repeat, as in a statistical experiment;’ another defines it as ‘to make or do something again in exactly the same way’ (*Merriam Webster Online Dictionary; Cambridge Advanced Learner’s Dictionary*). This clarity exists only at the level of dictionaries. Indeed, in a recent lawsuit involving economists the definition of replication constituted one of the two central issues, with the judge granting the defendant summary judgment by accepting the dictionary definition (*Lott v. Levitt* 2007). Here, I will refer to the dictionary definition as *pure replication*.

In economists’ common parlance the idea of replication clearly goes beyond this. A useful taxonomy was provided by the psychologist John Hunter (2001), who described *statistical replication* – different sample, but the identical model and underlying population. Statistical replication is only marginally relevant for us: I cannot imagine, for example, anyone taking a different sample from a major data set to examine the same model that has already been studied using those data. Presumably the original author used the entire data set, subject to excluding cases because of item non-response or disqualifying characteristics. One could, however, repeat the same lab experiment on a different set of undergraduate subjects, and that is the closest we can approach statistical replication. *Scientific replication* – different sample, different population, and perhaps similar but not identical model – appears much more suited in type to our methods of research and, indeed, comprises most of what economists view as replication.

In economic research, examining the same question and model using the underlying original data set amounts to pure replication (and in one of the controversies discussed below, Antonovics and Goldberger 2005, did exactly that). In the context of research by members of our profession, scientific replication also merits substantial attention. Using this taxonomy, I examine the history of, facts about, and prescriptions for pure replication in section 2. Section 3 focuses on scientific replication, particularly on internal replication, examining multiple sets of data within one study. Throughout I concentrate on the efforts that have been made to ensure replication and replicability, and I examine the incentives for replication facing the agents in the market for scholarly research in economics.

2. Pure replication

2.1. Recent history

Pure replication depends on the availability of all the information from the project that is to be replicated. Its potential importance was demonstrated to the profession in the mid-1980s by a project undertaken by the editor of the

Journal of Money, Credit and Banking (JCMB) (Dewald, Thursby, and Anderson 1986). He sought data sets and documentation from authors of recently published and accepted manuscripts from his journal, received them from a disturbingly small fraction (about one-third), and attempted to replicate the authors' results. In some ways the results were encouraging: Replication uncovered what the authors believed were many minor errors, but in only a few cases were the mistakes severe enough to alter the original work's qualitative conclusions. Of course, one wonders whether those data sets that were provided were not self-selected from the upper tail of quality of empirical work, so that the apparent relative lack of major errors underestimates the general severity of difficulties with published studies.

Even without positive selection, the results of the *JMCB* project should be disturbing to empirical researchers, to those theorists who pay attention to empirical work, and, most important, to the public generally, whose inchoate notions about economic phenomena are in the end partly formed by evidence generated through economic research. The findings clearly disturbed the editors of what is by far the most widely read refereed economics journal, the *American Economic Review (AER)*. Shortly after publishing the report on the *JCMB* experiment, the editors mandated that authors of accepted articles pledge to make their data sets available upon request unless the data were proprietary. Published articles never were required to state explicitly that the data were available, but a generalized editorial statement made that fact clear to the profession.

In response to some evidence that compliance with its policy was not satisfactory (McCullough and Vinod 2003), with the 2005 volume, the *AER* altered the policy to take advantage of technologically induced economies of scale by requiring authors of empirical papers to submit their data sets for inclusion on a special Web site maintained by the *Review*. Compliance with this requirement has been prompt and complete.¹ Authors can opt out of the requirement if the data set they used was proprietary and/or confidential. Regrettably, very few other journals (the *Journal of Applied Econometrics* and the *Review of Economic Studies* being exceptions) have yet adopted this technological advance. On its Web site this journal simply states, 'The objective is that a reader could, in principle, replicate exactly what was done in the paper,' but does not require submission of data sets for posting on the Internet. It does stress the need for ease of replication, however, and it does stipulate that authors should make it clear how they constructed their samples.

Other than the *AER*, none of the top-ranked general journals in economics explicitly requires provision of the data, and my perusal of the journals more generally suggests such statements are indeed quite rare. In my own subfield of labour economics, however, by the early 1990s two of the three leading publication

1 Report of the Managing Editor, American Economic Association, *Papers and Proceedings*, May 2006, table 8; May 2007, table 8. The editor of a specialized journal mentioned that he does not get involved in requests for data unless the author is slow to respond, which he notes has occurred only twice in five years.

TABLE 1
 Statistics describing data requests: articles in the *ILRR* 2002–4, *JHR* 2002–4, *AER* 1999–2000

Journal	Mean (std. error of mean)	Median	Response rate	<i>N</i>
<i>ILRR</i>	0.78 (0.24)	0	0.87	69
<i>JHR</i>	0.95 (0.21)	0	0.87	70
<i>AER</i>	8.88 (2.53)	3	0.83	60

outlets began requiring authors to include a statement about the availability of their data in the acknowledgment footnote (beginning with the 1990–1 volume of the *Industrial and Labor Relations Review* (*ILRR*) and with the 1991 volume of the *Journal of Human Resources* (*JHR*)). With the example of the *AER* confronting economic researchers, it is clear that authors today are expected to make their data available and that they can expect to field enquiries about their data.

With explicit statements of the availability of data and the development of the ethos that pure replication should be facilitated by authors of empirical studies, it is worth examining whether, in fact, these new tools for putting economic research on a more scientific basis have been taken up by economists. To do so I sent an email survey to authors of all 139 empirical studies published in the *ILRR* and *JHR* in 2002–4. Each author was asked how many times s/he had received requests for the data used in the published study (and whose availability was advertised on the first page of the article).

The first two rows of table 1 present the results of the survey of experiences at these two specialized journals. First, one should note that the response rate was unusually high. While unit non-response may be non-random, it is difficult to conjure up arguments why the non-respondents may have been self-selected from the upper tail of articles in terms of requests for data. Absent such arguments, I shall assume that the respondents are taken randomly from the distribution of empirical articles in these journals along the dimension of requests. The mean number of requests for data in each of these two specialized journals was just one. More important, 67% of authors of articles in the *ILRR* and 54% of those in the *JHR* never received a request for data. The histogram in figure 1 combines both journals' experiences. While most papers received no data requests, a few papers did receive a substantial number.² The conclusion from these data is that,

2 In statistical terms, the distribution is highly over-dispersed: A negative binomial approximates it far better than a simple Poisson distribution. One might worry that using examples from 2002–4 fails to allow enough time for requests to have been made. Indeed, time since publication for the *ILRR* is significantly positively related to the number of requests in a negative binomial regression; but for the *JHR* the opposite is surprisingly true.

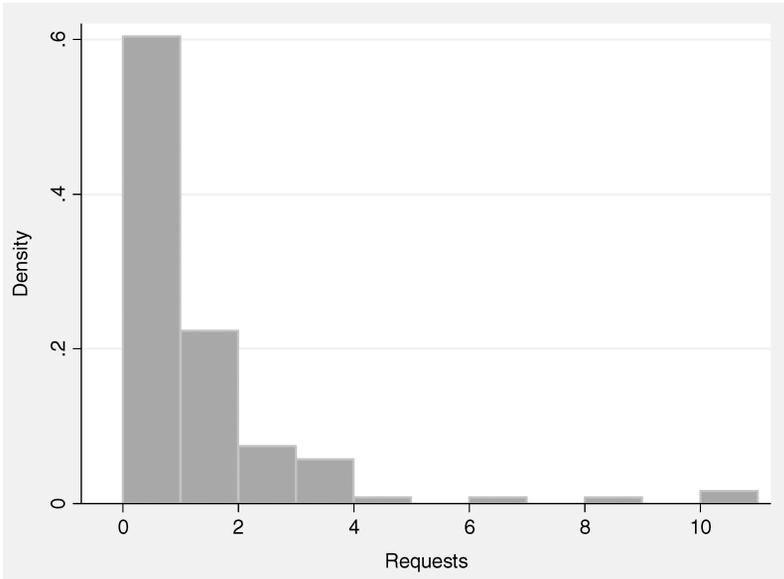


FIGURE 1 Distribution of requests for data, *Industrial and Labor Relations Review* and *Journal of Human Resources*, 2002–4

despite the effort to make data available, few researchers are availing themselves of the opportunities offered.

Perhaps the low demand for data and the small supply of replication studies suggested by the first two rows of table 1 arise from the relatively low readership of these journals and their specialized nature. To examine this possibility I sent the same email questionnaire to authors of empirical articles published in 1999 and 2000 in the *AER*, whose circulation is over three times that of any other refereed journal in our field.³ The response rate here, too, was excellent: 83%. A few of the responses may have been exaggerated (one author responded ‘over 100 times’), so that the mean of nine requests may be biased up; but data from the median article were requested three times, and only 22% of the authors received no requests for their data.

A fair conclusion from these little surveys is that, except for articles published in the most visible outlets, data sets are simply not requested for any purpose. One wonders whether even those that are requested are used for scholarly purposes or simply for pedagogy: One of my own data sets (Hamermesh and Biddle 1994) has been requested over 20 times, but all but one of these requests have been for use as an example in course work in econometrics or labour economics. Requests for classroom or textbook use are important for developing future scholars; the

³ Circulation of the *ILRR* is currently 2,400 and that of the *JHR* is 2,200. The circulation of the *AER* is over 21,000.

data sets are clearly being used, but hardly for pure replication for scientific purposes.

2.2. Incentives in the market for replication – positive aspects

One might consider the market for replication as containing three types of agents: editors of scholarly journals, who decide what to publish and what requirements about data to impose on authors of empirical studies; the authors of such studies; and other researchers, especially those who might potentially be interested in replicating a particular published work. Each set of agents has different goals, and changes in the technology of doing research and publishing alter their ability to meet those goals. In this subsection I examine for each agent how the possibility of replication might affect behaviour and how that interaction has changed with technology.

Journal editors presumably have one simple goal: Publishing articles that by their scholarly novelty and perhaps contemporary relevance maintain the current readership's interest and attract new readers. To achieve this goal the editor needs to make sure that the research is credible. This is less difficult for theoretical work: A properly chosen referee can check the lemmas and theorems and recommend rejection, or at least revision, if there are mistakes. No referee of empirical papers can or will re-estimate the models that the author has written up, so that the editor must rely on the author's credibility as a researcher and the general believability of the results. In almost all instances the editor is thus at the mercy of the authors and can only hope that the authors' incentives are compatible with the journal's.

Technological changes, in the form of greater portability of data and the frequent use of readily available, standardized data sets, have probably made the editor's job easier. Although such cases are very rare and less common than they should be, electronic communication does enable the editor to insist on verifying, or having an assistant verify, some particular calculation to satisfy his/her concerns about the submission.⁴ Even with lower costs of verification, however, it is unlikely that the editors will take the time or, in many cases, have the ability to verify the results of an empirical submission.

Every empirical paper that I have published over the last 40 years has involved repeated estimation and re-estimation of the underlying models. No study sprang Athena-like from my brow or that of my computer; all involved re-specifications, obtaining additional data to extend the model, and other extra work. Before the advent of the personal computer this meant keeping detailed chronological records of the computer output that I had generated – an econometric equivalent

4 One such occurrence involved a paper of mine (Hamermesh and Schmidt 2003), for which the journal editor insisted upon receiving the data set and presumably rechecked some of the estimates before agreeing to publish the piece. Initially I was extremely annoyed that the editor of this predominantly theoretical journal, himself a theorist, had the temerity to question my competence and/or honesty. After some thought, however, I was less bothered, wishing only that more people were subjected to this but wondering whether the editor might not have had something better to do with his time.

of a bench researcher's lab book. With a personal computer, the pile of output has been replaced by electronic files marked by date/time of production. Maintaining these records and the underlying data files has always taken time, and it would be much easier for me to be sloppy (or even sloppier than I am) in doing so. Nonetheless, the cost of keeping records has dropped over time.

The benefits of careful record-keeping and maintaining one's data files appear to have risen over time. As Ellison (2002) has shown, the number of rounds through which an economics paper goes before publication has increased over the past 30 years, and the time between initial submission to and eventual acceptance by the leading journals has more than doubled. That being the case, knowing that multiple revisions will be requested and recognizing that resurrecting a particular result from the distant, often multi-year, past may be difficult, empirical researchers today have greater incentives to maintain careful records.

Publishing a mistake has always left researchers open to general opprobrium from their fellows, and if the mistake is viewed as deliberate – as reflecting falsified data – can lead to the scholarly equivalent of Mennonite shunning (Kevles 1998), including the loss of funding and position. With instant gossip through email and blogs, the ease and speed with which one's mistakes might subject one to sanctions have increased. With any reasonable loss function – and most scholars are particularly risk averse along this dimension – changing technology has increased the benefits of careful documentation and maintenance of one's data sets for this reason, too.

For most scholars the payoff is in the influence of one's ideas: having other scholars base their work on those ideas, having students learn from them, and (an old man's perhaps vain hope) having public policy influenced by them (perhaps particularly important in the case of economics). The incentives we face here are clear: our ideas are unlikely to be taken seriously if our empirical research is not credible, so that the likelihood of positive payoffs to our research is enhanced if we maintain our data and records and ensure the possibility of replication. Technology has not altered these incentives; they have always been there. Here, too, however, the greater ease of communication worldwide may have enhanced these returns, particularly in the areas of influencing policy and stimulating students.

The third agent in the market for any empirical research is the worldwide set of researchers who might seek to replicate the study. For any individual the costs of replication are substantial: obtaining the data set and/or computer code and understanding what the author did. More important is the opportunity cost – with teaching and other duties, time for research is limited, so why spend it on replication if one has one's own ideas and data on which to test them? For most of us the return to spending time on replication is not great. First, most of the empirical studies are going to have at most only marginal impacts on economists' thinking. That is perhaps why we see such a difference between the experiences reflected in the first two rows of table 1 and those in the third row. Second, assuming authors are careful, the likelihood of the replication's providing substantial new knowledge rather than minor corrections is small. But the counter-example of

Feldstein (1974, 1982) and Leimer and Lesnoy (1982), probably the best-known replication issue in this profession in the past 40 years, suggests that this is not always so. As that example showed, and as theory would indicate, the likelihood of somebody attempting replication rises with the visibility of the published study and its author and decreases with the visibility of the potential replicating author. Under those circumstances the benefits of replicating are greater and the costs are lower.

Technology has diminished the costs of providing the materials necessary for replication at the same time that changes in the publication process in economics have increased the benefits to authors of maintaining the records that might make replication possible. Even without journals' requiring posting of data sets, or acknowledging the availability of data, I have no doubt that there would be greater possibilities for pure replication today than in the 1970s. So, were the requirements for making data available necessary when they were imposed? More important, are they still necessary?

I had a paper accepted by the *AER* shortly after it instituted its policy of accepting empirical articles conditional on the author acknowledging his/her willingness to make the data available. I was really angry about this new requirement, as I felt it increased the cost of doing empirical relative to theoretical work at a time when, in my view, too much useless theory was already being published. Why not impose a similar tax on theorists, for example, require submission of all the scribbled notes that led up to the work, or at least every last detail of proofs? Require theorists to submit photographs of blackboards on which they chalked down their ideas in front of colleagues! I calmed down after a bit, recognizing both that I could not do anything about the requirement and that the rigours it imposed on me might improve the quality of my work. Today, I have no doubt that this requirement and similar subsequent ones hastened the move towards an equilibrium in which empirical researchers behave in a way that is based on the expectation, or at least the fear, of having their work replicated.

I doubt that removing the current strictures that make replication easier would return us to the old equilibrium in which authors who potentially might replicate published work could not expect to obtain the data they desire. The cost of maintaining records is small. Also, I hope empirical researchers now recognize that it is in their own interest to maintain their data sets and make them available when asked. If nothing else, a few researchers' continuing willingness to make data available spills over onto others' behaviour, so that it is unlikely that removing requirements would shift the current equilibrium. Nonetheless, increasing the possibility of replication by maintaining these requirements or, better still, providing a centralized depository for data as in the new *AER* policy, seems especially desirable.

Does this availability mean that pure replication will take place? The survey evidence provided above suggests that in most cases replication is a threat that might keep potential cheaters honest rather than being a common practice. Indeed, two attempts to institutionalize replication work with which I am

familiar died from neglect. The *Journal of Political Economy* ran a section entitled 'Confirmations and Contradictions' from 1977 to 1999, occasionally publishing what were in most cases comments using new data, not pure replications. *Labour Economics*, which began publication in 1993, stated in its opening issue that it explicitly welcomed replication studies. After a few years, however, the invitation was dropped, despite the editors' commissioning a short piece (Hamermesh 1997) touting the virtues of replication. As the then editor wrote, there was a 'lack of interest: we simply got no submissions. There is a structural lack of interest in replication' (Hartog 2007). A similar lack of response is reported by the editor of the 'Replications' section of *Empirical Economics*.

An optimistic explanation for the apparent disinterest is that major errors (as opposed to the minor mistakes found by Dewald, Thursby, and Anderson 1986) are relatively infrequent. Editors have a clear bias in favour of publishing papers in which maintained hypotheses are refuted (DeLong and Lang 1992). Perhaps authors' interest in writing up the results of replications is spurred only when a central qualitative result of the original study is contradicted. And perhaps editorial decisions are motivated by the 'gotcha' mentality that pervades diverse aspects of modern life such as presidential politics and celebrity watching.

2.3. *Some normative conclusions based on replication controversies*

Economists cannot use the social-science equivalent of genetically identical laboratory rats. People do not behave that way, and, in any event, we are looking for social behaviour, not the purely biological responses of organisms. Nonetheless, making replication possible offers the social virtue of allowing those findings that seem most important to be verified or refuted directly on the particular data set used to generate them rather than examined later on a different and perhaps less appropriate set of data. If a finding is specious, better to have its props knocked out from under immediately than have it chipped away at slowly. This is especially important given how avidly the media pick up unusual findings (and how avidly some of us seek to publicize those findings, even highly preliminary ones, in the media). Given the incentives outlined above, I doubt that much pure replication will take place, but anything that aids it is praiseworthy.

To facilitate replication, if one receives a request for one's data sets or code, one should comply fully and speedily and should make the data available in a readily usable fashion. Among the few requests for data that I have made over the years, most were handled expeditiously and completely, but one was complied with in name only, the data being essentially useless to anyone except the original author.

Despite the paucity of replication studies, a number of comment-reply pairs involving at least in part pure replications have been published, and the controversies implicit in them have gripped the attention of many members of the profession (appealing, perhaps, to the same prurient interest as mud wrestling). I briefly review four recent examples, with the sole purpose of providing instruction

on how to behave in such situations, not to judge who was right (or even if there was a 'right' in the controversy). In reviewing them, one should remember that the replicating author usually views him/herself as on the attack. Usually, too, as noted above, the replicating author is more junior and sees the opportunity to make his/her reputation through discrediting a more senior economist's highly visible work. With these incentives it is all the more important to take a gentle, restrained, professional tone in the comment.

Leimer and Lesnoy (1982) made the bald statement, 'This paper presents new evidence that casts considerable doubt on Feldstein's conclusion.' While measured, but perhaps stronger than the work merited (although they did catch an important error), the statement is neither gratuitous nor *ad hominem*. In their long comment on Card and Krueger's (1994) evidence suggesting the absence of negative effects of higher state minimum wage rates on employment, Neumark and Wascher (2000) used both the original and new data, and concluded that 'the payroll data raise serious doubts about the conclusions CK drew from their data.' Antonovics and Goldberger (2005) engaged in pure replication to examine a study of the role of mother-child transfers of human capital (Behrman and Rosenzweig 2002) and concluded, 'We have seen that the results are not robust and that the policy inference may be misguided.' Again, this is a fairly mild conclusion, given the force of the arguments the authors had presented about problems with the underlying data.

Perhaps the best recent example of how to write a comment based on a replication is Easterly, Levine, and Roodman (2004), commenting on a recent paper (Burnside and Dollar 2000), whose data set was the third most frequently requested among those *AER* studies included in the sample reported in table 1. The comment extended the earlier data set to some additional countries and a few additional years, so that it falls somewhere between pure replication and scientific replication. Its main result was to demonstrate that the original finding – that the amount of foreign aid a developing nation receives interacts with good macroeconomic policy to induce growth but does nothing absent such policy – did not seem to be robust to the addition of relatively few data points.

If one complies with a request for data and shortly thereafter finds an article circulating that claims to have destroyed the principal findings of one's paper by challenging or at least casting considerable doubt upon them, what should one do? If a mistake was actually made, admit it honestly and immediately and move on to set out the importance of the error for your fundamental conclusions, thus 'limiting the damage' while also advancing closer to understanding the phenomena under consideration. The best model for this admission is Feldstein's (1982) 'Reply,' the first sentence of which was 'I am embarrassed by the programming error that Dean Leimer and Selig Lesnoy uncovered but grateful to them for the care with which they repeated my original study.'

Following in order the comments discussed above, consider the replies. Card and Krueger (2000) responded by analysing the new Neumark-Wascher data and obtaining additional data in a different manner from their original data

set (whose method of collection was the replicating authors' main concern) and concluded that 'the non-representative sample . . . produced [Neumark and Wascher's] anomalous results.' Behrman and Rosenzweig (2005) reworked the original data set and examined a new one, but concluded, using even stronger terminology than in the comment, 'We show that the focus of AG's comments is misguided, given the policy issues.' The Burnside and Dollar (2004) reply is a modern model of scholarship; but like the other replies it concluded that additional re-specifications lent stronger support to the original work.

By the time final versions of the comment and reply are in print and belong to the ages, replicating and original authors in these major controversies usually rise to appropriate standards of scientific discourse (whether voluntarily or at the insistence of an editor is unknown). When these significant scientific controversies play out in the press or in unpublished drafts, however, the comments and replies often are phrased in more strident tones. One example comes from a highly publicized (*Wall Street Journal*, 24 October 2005, 1) controversy about Hoxby (2000), on which, in an advance but as yet unpublished draft, Rothstein (2005) wrote, 'Hoxby has not provided the precise data set from which her published results were derived.' Hoxby (2005) countered and made a crucial point about the process of empirical research today: 'The original data simply do not exist and for a very good reason. I was trained to . . . [write] code that takes research . . . from the raw data to estimation. Such code may create intermediate datasets . . . but they are replaced every time the code is run . . . Obsolete datasets are not left sitting around to be used later, accidentally. The procedure prevents the unwitting propagation of erroneous or superseded data and code.'

In the four more recent controversies touched upon here questions of simple mistakes did not arise. Rather, the issues involved sample selection, choice of instrumental variables, questions about model specification, and other legitimate econometric concerns. In all but one case both new data were collected and the original data were reworked, so that most comments amounted to a mixture of pure and scientific replication. Readers of each should have been able to take away from the exchanges an assessment of each side's merit. If nothing else, the exchanges should have the additional benefit of offering students a better feel for the care needed in empirical work, particularly the importance of choice of sample and construction of variables. Both of these goals are enhanced if replicating and original authors temper their comments, not just in the published versions (as is true in all the examples discussed here) but also in the early drafts, so that they, and bystanders, including those who may be referees, can concentrate on the technical merit and the importance of the arguments.

One should remember that professional and public perceptions of facts do not in the end rest solely on the validity of one particular empirical study. Any alleged fact that affects the way that we think about a phenomenon will be hardened diamond-like by the heat and pressure of repeated empirical examination before it becomes background to our beliefs about how the world works. Nevertheless, given the media interest in reporting novel or titillating empirical findings and

politicians' desires to robe their proposals in scientific empirical cloth, however novel or inconsistent with prior research, it is crucial that as a profession we ensure that replication, or at least fear of replication, is our norm.⁵ Empirical economics is never going to become a laboratory science, but recognizing the role of replication can move us slightly in that direction by preventing us from propagating erroneous results.

2.4. *A modest proposal*

As the discussion in this section shows, replication is rare, mainly because the incentives for empirical researchers to perform replication studies are weak. Feldstein's (1982) prediction, 'replication studies . . . should become increasingly important,' has not proved to be accurate. Clearly, reductions in the cost of replication, and even the further reduction in cost created by central archives like that created by the *AER*, are unlikely to move the equilibrium quantity of replication studies very far from zero. Absent incentives on the supply side, what are needed are additional demand-side incentives for potential replicating authors. Given the reward structure in the economics profession, however, we cannot expect junior or even mid-level researchers to undertake replication studies. Even very senior economists, whose age-earnings and age-status profiles are at best flat, are unlikely to undertake replications without incentives beyond those that have been provided so far (editors' general expressions of interest in publishing such studies). A more proactive approach is required of editors of scholarly journals.

One arrangement consistent with the incentives I have discussed and that would generate additional replications would be for journal editors to commission leading senior empirical researchers to undertake a replication study of a paper of their choice, one that had previously been published in the journal. If editors of each of the three leading general journals commissioned two replication studies per year, with publication guaranteed subject to refereeing (NOT by the author of the original study) to assure some minimum quality level, more replications would be undertaken. Original authors would be expected to write a short reply to the final version of the replication study.

The cost of this proposal (the opportunity costs of the time of some very senior economists and of the journal pages that would otherwise be filled with what one hopes are the marginally acceptable papers) seems fairly small. The benefits, in terms of nudging the profession towards additional legitimacy and providing guidelines for researchers on how to conduct empirical research, seem substantial. Without this direct subsidy, however, I am certain that, while we will continue to see sporadic calls for more replication, the supply of replication studies will hover near zero.

In these studies the replicating author should expect reasonable cooperation from the original author. If the original author has provided documentation

5 The best example is President Clinton's seizing on the Card-Krueger results to justify a proposal to raise the federal minimum wage (State of the Union Message, 24 January 1995).

showing how s/he chose the samples and constructed the variables that underlay the published tabular material, s/he should not be required to do much more to help the replicating author. Nor should the replicating author expect much more: The original author cannot be expected to co-author a commissioned replication of his/her own work!

3. Scientific replications

3.1. *The sad state of scientific replication in economics*

Most of what economists view as replication represents scientific replication: re-examining an idea in some published research by studying it using a different data set chosen from a different population from that used in the original paper. The recent comments discussed in section 2 may have started as attempts at pure replication, but in most cases soon transformed themselves into hybrids of pure and scientific replication. There are a variety of reasons why scientific replication is likely to be more important in economics (and social sciences more generally) than in the natural sciences. With laboratory experiments one might argue that scientific replication is not very useful (but are the usual subjects of economic experiments – college students at selective American universities – credibly representative of the average consumer in Bourkina Faso?). The argument might be applied, albeit with less force, to field experiments, increasingly in vogue in the literatures in policy evaluation and mechanism design (see Levitt and List 2007). But for most empirical work we need to replicate analyses obtained on one data set many times over before the underlying and hopefully general economic point is to be believed.

By far the most important justification for scientific replication in non-experimental studies is that one cannot expect econometric results produced for one time period or for one economy to carry over to another. Temporal change in econometric structure may alter the size and even the sign of the effects being estimated, so that the hypotheses we are testing might fail to be refuted with data from another time. This alteration might occur because institutions change, because incentives that are not accounted for in the model change and are not separable from the behaviour on which the model focuses, or, crucially, that even without these changes the behaviour is dependent on random shocks specific to the period over which an economy is observed.

The same types of differences can generate non-robust estimates in studies based on one cross-section of data (from a particular area or economy). Institutions differ across and even within economies, and different random shocks buffet different cross-sections at any particular time. Even more of a problem arises from the nature of the social interactions that might generate, or at least affect the behaviour to which our hypotheses are addressed and for which we do not, and perhaps even cannot, account in our models. While social norms and ethnic capital have become an important focus of economic analysis in the past

TABLE 2

The distribution of empirical studies by type of problem and source of data used, recent volumes, *AER*, *JPE* and *QJE*, and *AJS/ASR*

	Type of issue							
	General			General (internat.)			Country-specific	
	US	Other	Many	US	Other	Many	US	Other
<i>AER</i> ($N = 115$) 2005–6	54	26	13	2	1	8	8	3
<i>JPE</i> ($N = 82$) 2004–6	49	12	10	0	1	3	5	2
<i>QJE</i> ($N = 95$) 2004–6	46	16	18	2	1	4	5	3
All three journals ($N = 292$)	149	54	41	4	3	15	18	8
<i>AJS, ASR</i> ($N = 122$) (2003–4)	75	14	18				7	8

15 years, their likely impact goes far beyond the few areas of behaviour in which they have been studied.

Underlying the importance of scientific replication is the dominance of American economists, and American data, in the study of economic phenomena. It is not the case that American journals have a bias towards publishing studies based on American data (Hamermesh 2002). Rather, researchers at American universities have for many years been pre-eminent in the profession, and U.S.-based journals have circulations that far exceed those of journals produced elsewhere. If our theories are intended to be general, to describe the behaviour of consumers, firms, or markets independent of the social or broader economic context, they should be tested using data from more than just one economy.

There is nothing wrong with our research jingoism, provided we make it clear that the statistical evidence we adduce about our hypotheses may well be applicable only to the United States, so that the verification of the idea is quite limited. One might argue that many such studies do not claim to be general – they deal with a specifically American problem – and that, in any case, U.S.-based journals should be dealing with specifically American problems. Unfortunately many studies do not stake out such a narrow area, and too often modesty about the broader applicability of U.S.-based results is missing.

Consider the results shown in table 2, which presents statistics describing studies published in regular issues of the *AER*, in 2005 and 2006, and in the *JPE* and *Quarterly Journal of Economics* (*QJE*) from 2004 to 2006, which use or are based on data. (This compilation thus includes the calibration literature.) I exclude comments, replies, presidential and Nobel lectures. I first divide the studies into

two groups, those that could not be interpreted (and that the authors do not interpret) as dealing with anything other than a country-specific issue and those that examine a general hypothesis or claim that the idea is generally applicable to economic behaviour. I further divide the latter group into studies that deal with international economics and those dealing with other topics. Finally, within each of these three categories I subdivide the studies into those based solely on American data, those based on some other single country's data (or on EU or Latin American data), and those that use data from at least two countries. It is only the last group of studies that demonstrates scientific replication within the same study, and those studies alone are ones that have any hope of vitiating concerns about the broader applicability of our empirical work.

That the country-specific studies are disproportionately based on U.S. data is unsurprising: after all, these are U.S.-based journals. Also unsurprising is the tremendous concentration of studies in international economics that use multi-country data: after all, their topic is *international* economics. The bulk of the 292 studies that are tabulated (85%) deal with general issues other than those in international economics, however, and of these 60% are based solely on U.S. data. Of the studies that deal with general issues outside international economics, only 16% use data from more than one country. Moreover, about half of those deal with issues in political economy where data from many countries are used to examine how institutional differences affect political and market outcomes. The conclusion from this brief examination of the pinnacle of the scholarly literature in economics is that it presents empirical research on general economic ideas that is to a great extent based on data representing one country, by far most often the United States.⁶

Are we economists particularly guilty of jingoistic testing of general ideas on American data? Not being knowledgeable about the categories/distinctions in other disciplines, I cannot answer this question broadly, and even a narrow answer is fraught with problems. To attempt an answer, however, I tried to classify two years of empirical articles in the *American Journal of Sociology* (*AJS*) (like the *JPE*, published by the University of Chicago Press) and the *American Sociological Review* (*ASR*) (the Association journal), the two leading journals in that field. I was unable to distinguish between general/international and other general research, so that I present the results for the two types combined. While categorizing studies in a field outside one's own is problematic, many of the topics and, to a lesser extent, the approaches in sociology are similar to those in parts of empirical economics.

6 While the research disproportionately uses the U.S. as the 'guinea pig,' a remarkable variety of other countries' data sets form the basis for the research published in these outlets. Data sets from 27 different countries, the EU, and Latin America underlie the empirical research published recently in these journals. As a bibliometric note, a test of the hypothesis that the patterns of publication by type across the three journals are independent cannot reject that hypothesis ($\chi^2(14) = 9.28$ – not significantly different from zero at any conventional level). Implicitly, these three general journals publish remarkably similar types of research based on remarkably similar types of data (classified by country of origin).

I present the results of this little exercise in the bottom row of table 2. The distribution looks remarkably like that for the aggregate of economics journals.⁷ More than half of all the papers can be classified as claiming general implications based on empirical analyses performed solely on U.S. data. The classifications suggest that economists are no more guilty of intellectual ethnocentrism than practitioners of this related social science.

Now there may be reasons for immodesty and for the absence of much within-study scientific replication. First, and most obviously, the profession puts a premium on the creativity and generality of the idea, not on verifying the breadth of its applicability. Also, the academic world offers us a variety of incentives to generate a coterie of followers of our research ideas. Engaging in scientific replication within an original study reduces the opportunity for others to publish scientific replications of our work and thus reduces the potential size of a scholar's group of disciples. One might argue that many of the publications in second-level general journals and in many specialized journals too are essentially scientific replications (on similar data sets for a different country or different data sets for the same country) of a publication in a major general journal. Those publications enhance the original scholar's professional standing, and we can maximize them by producing a novel and broadly applicable idea that seems supported by the one data set we used but that leaves wide room for additional confirmation/contradiction/extension.

I doubt that many (any?) empirical researchers consciously engage in such gaming. The rarity of scientific replication within studies may well arise from the extra data work it would involve. In order to perform a within-study replication one must become familiar with a different data set, often one that is not coded in English. It is also the case, however, that the incentives for doing within-study scientific replication are non-existent.

One might dismiss this concern by saying that adding a second sampling base to one's study will merely increase the reader's confidence in one's findings (assuming the results are similar) in proportion to $\sqrt{2}$. Whether one views it approvingly as Baconian scientific method or disapprovingly as specification search (Leamer 1983), testing an idea on a single set of data is likely to lead to the conclusion of most interest to the author(s) (and to journal editors). A second data set, one that does not merely sample from the same population as the first, should enhance a reader's confidence in the validity of the results by a proportion far above $\sqrt{2}$.

3.2. *What is to be done?*

Given the incentives in the academic market, the outcomes suggested by the compilation reported in table 2 are unsurprising: The rewards to within-study scientific replication are small. The compilation does not clarify whether editors

⁷ When the distribution of all economic studies, except those classified as general-international, is compared with that of the sociological studies, a test that the distributions are identical cannot reject the null hypothesis ($\chi^2(4) = 6.89$).

do not care much about within-study scientific replication or whether the supply of such studies is tiny. I should imagine, however, that editorial demand will create its own supply, so that the crucial issue is altering editorial recognition of the importance of scientific replication. No editor of a major journal is likely to publish replications of previous original pieces. Also, why should an editor wish to burden empirical researchers with more work, and why should the relative burden compared with writing purely theoretical papers be weighted still further?

The answer must be the unappealing equivalent of the moral imperative: because it is right for the profession. The best empirical economists are quite adept at writing down clever, novel models that are wonderfully consistent internally, that are intellectually beautiful, and that are supported by the single set of data on which they are 'tested.' As I noted above, however, the validity of these tests is questionable, and adding a second data set will enhance their validity more than proportionately. The most credible studies I have done are those that have proposed an idea (hopefully a fairly novel one) and then tested it either on data sets from different economies or at least from the U.S. at different points in time. Policy makers and the general public will take the work that editors choose to publish (and the unpublished studies that, given editorial lags, circulate widely pre-publication) much more seriously if they do include scientific replications.

One cannot expect editors to require all authors to include within-study scientific replications. In some cases data sets are unique to a particular idea, even though the model may be universally applicable. In other cases, however, the data set is just one of many equally useful and generally available sources. For example, why the concentration on the American PSID when the equally readily available and quite comparable Australian HILDA, British BHPS, Canadian SLID, and German SOEP are readily at hand? Editors and referees should be sufficiently aware of these cases to require authors of papers that appear acceptable for publication to broaden their empirical focus beyond the single (usually American) data set they have used.

One might take the Alfred E. Neuman approach to scientific replication: 'What, me worry?' about studies based on just one set of data. Eventually, an idea will be tested in the market by subsequent empirical studies that are usually published further down the food chain of journals. Also, survey articles and meta-analyses of the leading and following studies can eventually tease out general facts about the particular phenomena (assuming that such facts exist).

One difficulty with this Panglossian view is that these follow-ups are usually not brought to fruition until many years after the initial original empirical piece was published. During that hiatus perceptions of the truth can easily crystallize around results that are not generalizable and perhaps not even essentially correct. Also, academic incentives are such that the people undertaking the surveys and meta-analyses do not usually have the same professional visibility or credibility as the author of the original study (and may not be viewed as being as competent professionally). For these reasons it is crucial that editors of the leading

journals tilt the publishing process a bit more in favour of within-study scientific replication.

4. Conclusions

The cost of pure and scientific replication in economic research has diminished over the past 40 years. At the same time our verbal nods towards the need for replication have increased (as shown by the authors' common practice, in the replies discussed in section 2, of thanking the authors of the comments and stressing the great scientific benefits that result from replication). Despite the declining costs and despite our expressions of preferences, both pure and scientific replications are very rare in leading journals, and even within-study scientific replication is unusual.

Our expressions of preference are cheap talk. The profession provides few incentives for most active economists to produce replications of others' research, and similarly few to increase the believability of one's own research by testing ideas on multiple sets of data. Any demand for replication must arise from the profession as a whole, as intermediated by the actions and decisions of the editors of the most visible professional journals. Editors need to take the lead by providing sufficient incentives for top-flight authors of empirical work to engage in pure and scientific replication and by insisting on within-study scientific replication in many more articles. Without these changes occasional paeans to the virtues of replication are as likely to enhance the scientific soundness of empirical research in economics as programs that urge abstinence are to reduce teenagers' sexual activity.

References

- Antonovics, Kate, and Arthur Goldberger (2005) 'Does increasing women's schooling raise the schooling of the next generation? Comment,' *American Economic Review* 95, 1738–44
- Behrman, Jere, and Mark Rosenzweig (2002) 'Does increasing women's schooling raise the schooling of the next generation?' *American Economic Review* 92, 323–34
- (2005) 'Does increasing women's schooling raise the schooling of the next generation? Reply,' *American Economic Review* 95, 1745–51
- Bernanke, Ben (2004) 'Editorial statement,' *American Economic Review* 94, 404
- Burnside, Craig, and David Dollar (2000) 'Aid, policies and growth,' *American Economic Review* 90, 847–68
- (2004) 'Aid, policies and growth: reply,' *American Economic Review* 94, 781–4
- Card, David, and Alan Krueger (1994) 'Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania,' *American Economic Review* 84, 772–93
- (2000) 'Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania: reply,' *American Economic Review* 90, 1397–420

- De Long, J. Bradford, and Kevin Lang (1992) 'Are all economic hypotheses false?' *Journal of Political Economy* 100, 1257–72
- Dewald, William, Jerry Thursby, and Richard Anderson (1986) 'Replication in empirical economics: the *Journal of Money, Credit and Banking* Project,' *American Economic Review* 76, 587–603
- Easterly, William, Ross Levine, and David Roodman (2004) 'Aid, policies and growth: comment,' *American Economic Review* 94, 774–80
- Ellison, Glenn (2002) 'The slowdown in the economics publishing process,' *Journal of Political Economy* 110, 990–1034
- Feldstein, Martin (1974) 'Social security, induced retirement and aggregate capital accumulation,' *Journal of Political Economy* 82, 905–26
- (1982) 'Social security and private saving: reply,' *Journal of Political Economy* 90, 630–42
- Hamermesh, Daniel (1997) 'Some thoughts on replications and reviews,' *Labour Economics* 4, 107–9
- (2002) 'International labor economics,' *Journal of Labor Economics* 20, 709–32
- Hamermesh, Daniel, and Jeff Biddle (1994) 'Beauty and the labor market,' *American Economic Review* 84, 1174–94
- Hamermesh, Daniel, and Peter Schmidt (2003) 'The determinants of Econometric Society fellows elections,' *Econometrica* 71, 399–407
- Hartog, Joop (2007) Email communication, 2 February
- Hoxby, Caroline (2000) 'Does competition among public schools benefit students and taxpayers?' *American Economic Review* 90, 1209–38
- (2005) 'Does competition among public schools benefit students and taxpayers? A reply to Rothstein,' Working Paper No. 11216, National Bureau of Economic Research
- Hunter, John (2001) 'The desperate need for replications,' *Journal of Consumer Research* 28, 149–58
- Kevles, Daniel (1998) *The Baltimore Case* (New York: Norton)
- Leamer, Edward (1983) 'Let's take the 'con' out of econometrics,' *American Economic Review* 73, 31–43
- Leimer, Dean, and Selig Lesnoy (1982) 'Social security and private saving: new time-series evidence,' *Journal of Political Economy* 90, 606–29
- Levitt, Steven, and John List (2007) 'Viewpoint: On the generalizability of lab behaviour to the field,' *Canadian Journal of Economics* 40, 347–70
- Lott v. Levitt* (2007) Northern District of Illinois, ruling of 11 January, Judge Ruben Castillo
- McCullough, Bruce, and H.D. Vinod (2003) 'Verifying the solution from a nonlinear solver: a case study,' *American Economic Review* 93, 873–92
- Neumark, David, and William Wascher (2000) 'Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania: comment,' *American Economic Review* 90, 1362–96
- Rothstein, Jesse (2005) 'Does competition among public schools benefit students and taxpayers? Comment,' Working Paper No. 11215, National Bureau of Economic Research